The Impact of Uncertainty Shocks on the Job-Finding Rate and Separation Rate

Markus Riegler*

April 5, 2018

Abstract

This paper shows that higher uncertainty increases unemployment in the US by increasing the separation rate *and* decreasing the job-finding rate. Standard search and matching models predict an *in*crease in the job-finding rate instead. I develop a search and matching model in which heterogeneous firms are subject to decreasing returns to scale and can hire multiple workers. In this framework, job flows (job creation and job destruction) do not necessarily coincide with worker flows (hires and separations). Costly job creation is key to obtaining a decrease in the job-finding rate after an increase in uncertainty.

1 Introduction

It has been suggested that an increase in uncertainty during the Great Recession has contributed to higher unemployment in the United States. There are two channels through which uncertainty shocks can affect the unemployment rate. First, it is possible that higher uncertainty increases the job separation rate so that employed workers are more likely to lose their jobs. Second, higher uncertainty could reduce the job-finding rate, which makes it harder for the unemployed to find a job.

In my empirical analysis, I confirm that higher uncertainty reduces employment. This reduction is driven by both a *higher separation rate* of workers *and* a *lower job-finding rate* of unemployed. The reduced job-finding rate contributes more to higher unemployment than the separation rate after an increase in uncertainty. Existing search and matching models of the labour market cannot account for this significant contribution of the job-finding rate.

In the literature, there are several measures of uncertainty both on the micro-level and on the macro-level. In the context of this paper, higher uncertainty means a higher expected *and* higher realized dispersion of idiosyncratic productivity across firms. Bloom et al. (2018) use microdata to show that plant-level shocks become more dispersed in recessions. This makes recessions, and in particular the Great Recession, appear to be a combination of a negative first-moment shock and a positive second-moment shock at the establishment-level.

An increase in dispersion of productivity across firms implies that more firms get hit by large negative shocks that lead to the layoff of workers. This so-called *realized volatility effect* raises the separation rate in times of higher uncertainty. At the same time it also means that more firms experience large positive shocks, which can increase hiring. Therefore, a reduction of the job-finding rate as seen in the data cannot be explained by the realized volatility effect. In contrast, what is known as the *wait-and-see effect* has the potential to reduce the job-finding rate: if firms are more uncertain about future realizations of productivity, they might become reluctant to hire workers in

the presence of adjustment costs. This happens if higher uncertainty reduces the expected continuation value of a newly hired worker.[1] For a model to explain the fall of the job-finding rate, it is important that a wait-and-see effect that makes firms more reluctant to hire exists and that it is sufficiently strong relative to the realized volatility effect.

The standard search and matching model assumes constant returns to scale, which rules out meaningful heterogeneity in firm size. If there are no idiosyncratic shocks, firm size is irrelevant and not pinned down. If there are idiosyncratic productivity shocks, as in Mortensen and Pissarides (1994), firms cannot be allowed to employ an unrestricted amount of workers. Because if it were possible to hire more workers at a given productivity level, the most productive firm would hire all workers. An implication of this restriction to hire additional workers is that the value of a matched worker to a firm increases without bounds with idiosyncratic productivity. In contrast, if the firm experiences lower productivity in a future period, it has the option to destroy the match. Therefore, the lower bound for the value of the match is 0. This asymmetry turns out to be important for the effect of uncertainty on labour market variables. If volatility of idiosyncratic productivity increases, more matches are destroyed by sufficiently negative realizations of productivity. Therefore, the realized volatility effect increases the separation rate. In addition, the absence of an upper bound implies that the expected value of a match increases. Hence the wait-and-see effect makes firms more reluctant to fire but it actually *increases* the expected value of a new match. This last effect increases vacancy posting and the job-finding rate in response to higher volatility, which is at odds with patterns observed in the data.

I develop a more general framework that does allow for meaningful heterogeneity across firms. In particular, I assume that a firm's production function exhibits decreasing returns to scale, which means that there can be idiosyncratic productivity shocks without restricting firms in the number of workers

---

[1]By the same logic, firms become more reluctant to fire workers if the expected continuation value increases in times of higher uncertainty because large improvements of productivity become more likely. The firm can wait for such a shock and fire the worker in the future if productivity deteriorates.

they can hire. In equilibrium, firms with various levels of idiosyncratic productivity will hire workers if they are small relative to their productivity, because their marginal product is high. To be precise, firms hire workers until the marginal value of an additional worker equals the hiring cost.

In this model, it is possible for the wait-and-see effect to reduce hirings if higher uncertainty reduces the marginal value of a worker. The marginal value of a worker consists of the marginal profit in the current period plus the expected continuation value the firm gets from the worker being at the firm in the next period. To understand the expected continuation value of a match it is important to consider the contribution of this one worker in future periods for different realizations of the idiosyncratic productivity shocks. Imagine a firm becomes very productive in the next period. In the standard framework, this means that the match becomes very valuable for the firm. In a framework in which the firm is allowed to hire more than one worker, however, the marginal value of one additional worker can never exceed the hiring cost of an additional worker. In this case, the continuation value of a worker hired in the current period will be equal to the discounted hiring cost that the firm saves in the next period. The ability to hire additional workers, when productivity is high, is of course valuable to the firm. But what matters for a firm's optimal choice is the value of the marginal match and this value is bounded by the option to hire additional workers. This puts an upper bound on the continuation value of a worker when the firm becomes highly productive.

The upper bound of the marginal value function is a key difference compared to the standard framework. It is an implication of the option to hire additional workers, comparable to the lower bound of the continuation value implied by the option of destroying matches. The upper bound introduces a concave part to the marginal value function, which has important implications for the question of how uncertainty affects the job-finding rate. Due to the concavity, it is possible that higher idiosyncratic volatility reduces the value of the marginal worker. In turn, this reduces the number of vacancies posted by the firm and the job-finding rate for the unemployed.

In my model, I distinguish between workers and the positions they fill.

There are sunk costs of creating a position. They do not have to be paid again when a firm wants to replace a worker that leaves. For example, if a worker turns out to be not a good match for the firm, or if he finds a job at another firm, his workplace still exists at the previous firm. It is then less costly for the firm to hire a new worker who does the same task. The creation cost could be interpreted as the share of capital needed for a position that cannot be recovered if it is shut down, or as the cost of finding new clients when a firm increases its size. As a result of this sunk cost, some firms in my model will fill existing empty positions even though they do not want to create new positions. These firms shrink over time, because some of their positions become obsolete. But they still hire workers, because the quit rate of existing workers exceeds the rate of obsolescence. This is consistent with differences between worker and job flows as documented by Davis et al. (2012): using JOLTS establishment data, they calculate that hires amount to around 10% of employment at shrinking establishments on a quarterly basis.

I show that the creation cost is important for the response of the job-finding rate to uncertainty shocks, because it strengthens the wait-and-see effect: as explained above, the marginal value of a position is bounded from above by the cost of creation, because the firm always has the option to create another position. The marginal value is at the upper bound, if a firm creates positions in the current period. Therefore, it does not increase if that firm becomes more productive in the future. In contrast, negative productivity shocks reduce the marginal value of a position. Higher uncertainty thus reduces the expected marginal value, making firms more reluctant to create positions and hire workers.

In order for the wait-and-see effect to be quantitatively important, it is necessary that interior continuation values are sufficiently likely.[2]

Due to the cost of creating positions in my framework, a firm becomes less likely to destroy positions after a negative shock. It will rather keep the

---

[2]If the continuation value was at its uppper bound after each positive productivity shock and at its lower bound after each negative productivity shock, increasing the variance of the productivity shock would not affect the expected value. The reason is that the continuation value would neither be strictly concave nor convex but flat at the realized values.

existing positions and workers, and it might even want to fill empty positions when workers quit. In these cases, the marginal value of a position takes an interior value. Then, higher uncertainty can have significant negative effects on the expected marginal value of a position created today, and this strong wait-and-see effect can reduce the number of posted vacancies. As a result, the job-finding rate falls as observed in the data.

The wait-and-see effect also becomes stronger for those firms that want to fire workers. This means that they reduce the number of workers that they fire which dampens the increase of the separation rate due to the realized volatility effect. As a result, the separation rate contributes less to changes in unemployment in accordance with the data.

In the absence of the creation cost, the job-finding rate hardly reacts to changes in uncertainty. In this case, higher uncertainty increases unemployment only because of a higher separation rate. When the cost is sufficiently large, however, the model can generate a significant drop of the job-finding rate in times of higher uncertainty. Then, my model is not only able to explain the changes of the unemployment rate in the last decade, but it can also account for the steep drop in the job-finding rate during the Great Recession.

My model is related to the literature on multi-worker firm models with decreasing returns to scale and idiosyncratic productivity shocks. Examples of random search models include Cooper et al. (2007), Elsby and Michaels (2013), and Fujita and Nakajima (2016), whereas Schaal (2017) and Kaas and Kircher (2015) assume directed search. The difference in my model is that it is costly for firms to create positions, which is an important assumption to explain the falling job-finding rate in times of higher uncertainty. With the exception of Schaal (2017), these papers do not consider uncertainty shocks. Schaal (2017) shows that it is difficult to explain the changes of the unemployment rate during the Great Recession while matching observed labour productivity, when only aggregate productivity shocks are used. His model with uncertainty shocks can explain the increase of the unemployment rate better. But it cannot explain the reduction of the job-finding rate during the Great Recession as the increase of the unemployment rate is primarily driven by a higher separation

6

rate. Because of the creation cost present in my framework, the job-finding rate falls in response to higher uncertainty, which fits the observed data better.

The paper is structured as follows: the next section provides evidence on the effects of uncertainty shocks on the job-finding rate and the separation rate in the US. Section 3 describes the model and derives the optimal choice of firms. Section 4 describes the calibration that is used. In section 5, I show the effects of uncertainty shocks in my model. I analyze how good the model can fit the US data from 1998-2013 in section 6 and section 7 concludes.

2 Effects of uncertainty in US data

In this section, I construct a measure for uncertainty based on the implied volatility of US stock options. I then use it to estimate impulse response functions of labour market transition rates following uncertainty shocks. In the theoretical model, uncertainty refers to the second moment of productivity on the firm level. Therefore, I want to construct a measure based on the volatility of *individual* firms as opposed to the volatility of an index of firms.[3,4]

I use data on implied volatility of US stocks, available from 1996 to the second quarter of 2013. For each firm, I calculate the quarterly average of log implied volatility of its call options with a maturity of 30 days. Allowing for firm fixed effects, I estimate time dummy variables for each quarter. The resulting time-series is shown in Figure 1.

The constructed measure for uncertainty is used to estimate a VAR including volatility, labour productivity, and quarterly averages of the job-finding rate and separation rate. The following Cholesky ordering is used to identify volatility shocks: I assume that labour productivity does not react instantaneously to volatility shocks and shocks to the labour market transition rates.

---

[3]For example, the VIX measures the implied volatility of the S&P 500 index. To the extent that idiosyncratic firm shocks are uncorrelated, higher volatility at the firm level does not directly translate into higher volatility of the index because of diversification.

[4]Leahy and Whited (1996) and Bloom et al. (2007) use the volatility of stock returns as a measure of uncertainty. An advantage of using the implied volatility of options is that it is a forward-looking measure, but it is not available as far back in time.
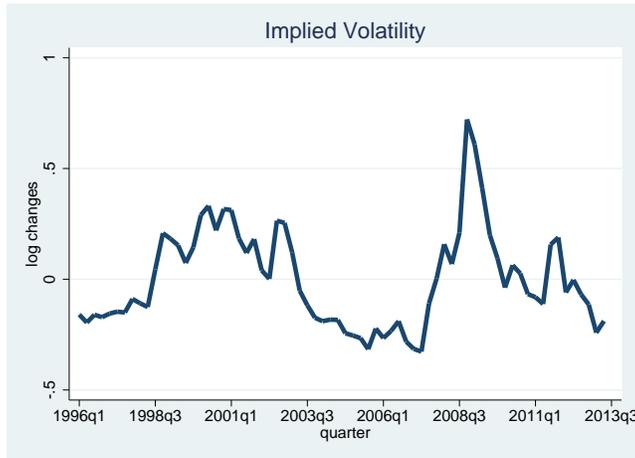
Figure 1: Index of implied volatility constructed from US stock options

Furthermore, volatility may only react to shocks in the transition rates with a one period lag.[5]
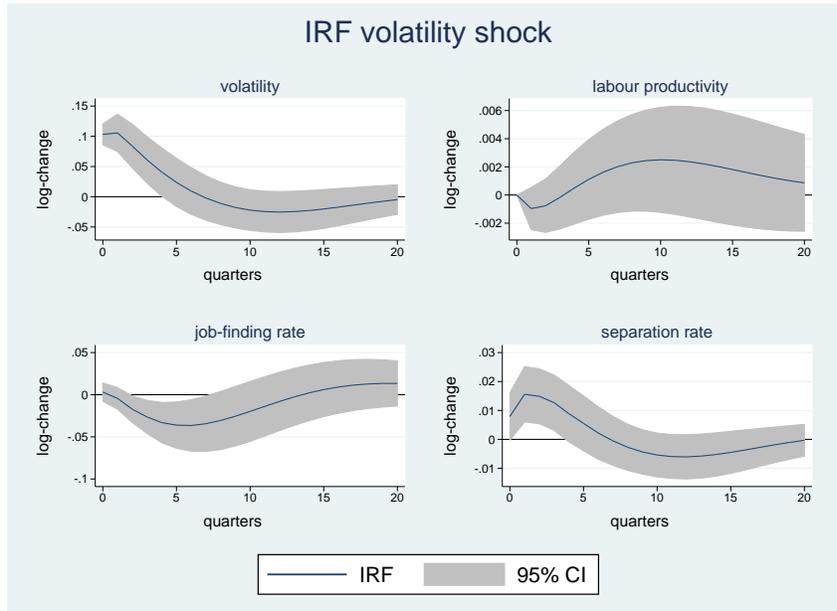
Figure 2 shows the impulse response functions after a one standard deviation shock to volatility. The job-finding rate falls and the separation rate increases. Both effects contribute to an increase in unemployment following an increase in volatility. Note that the response of the separation rate peaks earlier than the job-finding rate, whose fall is more persistent. Its relative change is initially stronger but after a couple of quarters the job-finding rate changes more. These relative changes determine how much the unemployment rate is affected by the shock. As shown in Appendix A.2, the following approximation holds:

$$\log \frac{u}{\bar{u}} \approx (1 - \bar{u}) \left(\log \frac{s}{\bar{s}} - \log \frac{f}{\bar{f}} \right). \tag{1}$$

This means that a 1% increase of the separation rate affects the unemployment rate by as much as a 1% decrease of the job-finding rate. Intuitively, this holds because flows into and out of unemployment are equal on average.

The negative effect of higher uncertainty on the job-finding rate that I find is in line with other results in the literature. Recent empirical studies by

---

[5]Appendix A.3 shows that the results are robust to changing the Cholesky order, as well as using a different measure for uncertainty, and different detrending of the data.

Figure 2: Impulse response functions after a volatility shock

Guglielminetti (2016) and Mecikovsky and Meier (2017) use different measures for uncertainty and find similar effects. Guglielminetti (2016) uses a survey-based measure of uncertainty. In a trivariate VAR, she finds that the job-finding rate falls after an increase in uncertainty. Mecikovsky and Meier (2017) use the measure of macroeconomic uncertainty in the US proposed by Jurado et al. (2015). They find that job creation falls, wheras job destruction increases in response to higher uncertainty.

3 Model

This section describes the framework and derives the equations that are needed to solve the model.

3.1 Setup

There is a constant unit mass of workers, who search for jobs when they are unemployed and earn wages when they are employed. They are assumed to

be risk-neutral. Workers can lose their job for two reasons. First, they can be fired when their employer experiences a sufficiently negative productivity shock that makes the match unprofitable. Second, it is assumed that not all separations are due to low productivity of the job. With a certain exogenous probability workers quit their job at the end of each period. In both cases, workers that become unemployed enter the unemployment pool in the next period, and can potentially find a new job.

I consider a constant mass of firms that can each hire multiple workers and face idiosyncratic productivity shocks.[6] Firms post vacancies and matching takes place according to a constant returns to scale matching function. The matching function can be fully characterised by the probability $q\left(\theta\right)$ that a vacancy is filled, where labour market tightness θ is the ratio $\frac{v}{u}$ of posted vacancies and unemployed workers. There are two types of costs that a firm incurs when it wants to increase its size. First, it has to pay a cost H for each position that it creates. Second, it has to pay a cost c for each vacancy it posts in order to fill empty positions. The firm takes labour market tightness and the implied probability of filling each vacancy as given.

In each period, a share δ of positions becomes obsolete. In this case also the workers who filled the positions, leave the firm. In addition, a share λ of workers quit without rendering the respective positions obsolete. In this case, the firm can hire another worker for a now vacant position and it only incurs the hiring cost $\frac{c}{q(\theta)}$ but not the creation cost H.[7] A firm can also endogenously fire workers and shut down their positions without any cost, if

---

[6]There is no firm entry and exit in my model. This simplification can be justified by the low cyclicality of job creation at start-ups as documented by Coles and Kelishomi (2011). Similarly, Fujita and Nakajima (2016) argue that cyclical fluctuations of job flows are mostly accounted for by the expansion or contraction of existing establishments.

[7]Fujita and Ramey (2007) also assume a one-off cost of creating positions, which has to be paid before posting vacancies. After separations that are not due to obsolescence, empty positions can be filled again. The important difference with my framework is that firms are not subject to idiosyncratic productivity shocks and operate at constant returns to scale in Fujita and Ramey (2007). As a result, the value of each empty position always equals its cost of creation, which makes the model very tractable. But the lack of meaningful firm heterogeneity would not allow to analyze the effects of increases in idiosyncratic volatility. In contrast to Fujita and Ramey (2007), I do not assume that the cost of creating positions increases in the total number of positions created.

it wants to reduce its size.[8]

For simplicity, I assume that a law of large numbers holds when the firm posts vacancies. That is, the firm can fill each empty position for sure by posting $\frac{1}{q(\theta)}$ vacancies. Therefore, the cost of filling an empty position is $\frac{c}{q(\theta)}$.[9] Firms are not allowed to keep an empty position idle. If it does not want to fill an empty position, it has to close it, and pay the cost H again if it wants to reopen it in the future. Even though firms cannot mothball positions, the distinction between positions and workers makes the model more interesting. Now each firm has to decide at the beginning of the period whether it wants to fill its empty positions. Firms that are relatively productive will decide to fill these positions whereas other firms will close at least some of them. As a result, some firms hire workers, while not creating new positions or even shrinking due to obsolescence. This is in line with the empirical results of Davis et al. (2012). They find for shrinking establishments that hires amount to around 10% of employment on a quarterly basis.

The production function takes the form $zxF(n)$, where z is the aggregate productivity and x the idiosyncratic productivity component. Both follow a Markov process. The total wage bill is denoted by W. The wage equation is specified in equation (15). In general, it depends on both firm-specific and aggregate variables. The total wage bill is increasing in the number of workers

---

[8]I assume throughout the paper that a cost is incurred when a position is created. It is a minor change to the model to let firms pay a cost when it closes positions instead. The results when a destruction cost is used instead of the creation cost are very similar for the following reason. When a firm opens a position, it knows that with certain probabilities it will destroy it in future periods when it becomes sufficiently unproductive. Therefore, the firm takes into account the net present value of this future destruction cost upon creating positions. The resulting optimal creation of positions becomes similar even though the firm only incurs the cost when it destroys the position. Figure B.1 in Appendix B.1 shows the results of the comparative statics exercise when a destruction cost is used.

[9]In an earlier version of this paper, I assumed that vacant positions only get filled gradually. This means that both the number of filled and the number of empty positions become state variables of each firm. Adding another state variable makes the numerical solution of the firm's problem more costly. The impact of the simplifying assumption on the results is likely to be small, because it does not take firms long to fill vacancies on average. For example, den Haan et al. (2000) use a quarterly job-filling rate of 0.71, which is in line with the filling probability found by van Ours and Ridder (1992) for establishments in the Netherlands.

and in firm productivity. In addition, it is increasing in aggregate productivity and market tightness. These are typical properties of wage equations in search and matching models, when workers and firms bargain over the surplus of their match. First, workers get a higher wage when the firm is more productive. Second, they get a higher wage when market tightness is high, as it would be easier for a worker to find a new job when he becomes unemployed, whereas it is more difficult for the firm to replace the worker.

The timing in each period is as follows: at the beginning of the period, a firm with n_{-1} positions has $(1 - \lambda) n_{-1}$ workers that fill them. The remaining λn_{-1} positions are empty. Then, the shocks of the aggregate and idiosyncratic productivity processes materialise. Afterwards, the firm has to decide, whether it wants to create new positions. In addition, it has to decide how many of its empty positions it wants to fill.[10] The firm shuts down all vacant positions that it does not want to fill. In the next step, the matching process takes place. Afterwards, the firm can fire workers and close their positions.[11] Then production takes place. The number of workers in the production phase is denoted by n^*. After production, a fraction δ of positions becomes obsolete and the respective workers become unemployed. Finally, a fraction λ of workers quits. Hence, the firm begins the next period with $(1 - \delta) n^*$ positions, filled with $(1 - \lambda) (1 - \delta) n^*$ workers.

3.2   Solution to firm problem

Let s denote the aggregate state. It summarizes the current realization of the aggregate shocks as well as the distribution of positions and idiosyncratic productivity across firms. $\Pi(n, x, s) := zxF(n) - W(n, x, s)$ denotes the firm's revenue minus the total wage bill. The value of a firm with n_{-1} positions at the beginning of the period and idiosyncratic productivity x is denoted by

---

[10]It is always optimal for a firm that creates new positions to also fill its empty positions.

[11]Firing takes place after the matching process such that only workers who were unemployed at the beginning of the period can be matched in that period. When I calibrate the model, one period corresponds to one week, such that this assumption is quantitatively not important.

$V\left(n_{-1}, x, s\right)$, when the aggregate state is s. It can be written recursively as[12]

$$V\left(n_{-1}, x, s\right) = \max_{n} \Pi\left(n, x, s\right) - \left(n - \left(1 - \lambda\right) n_{-1}\right)_{+} \frac{c}{q\left(\theta\left(s\right)\right)} - \left(n - n_{-1}\right)_{+} H$$
$$+ \beta \mathbb{E}\left[V\left(\left(1 - \delta\right) n, x', s'\right) | x, s\right]. \quad (2)$$

The firm's per period profit is given by the difference between production and the total wage bill reduced by the cost of filling $\left(n - \left(1 - \lambda\right) n_{-1}\right)_{+}$ positions and the cost of creating $\left(n - n_{-1}\right)_{+}$ positions. The continuation value is given by the expected discounted value of a firm with $\left(1 - \delta\right) n$ positions at the beginning of next period.

Note that there are two kinks in the objective function that the firm wants to maximise. The first kink appears when $n = \left(1 - \lambda\right) n_{-1}$. Then, the firm does not fill any of its empty positions, but it also does not fire any of its existing workers. The second kink occurs when $n = n_{-1}$. In this case, the size of the firm stays exactly the same within the period. It fills all its existing empty positions, but it does not create any new positions. The firm's optimal choice could be either at one of these two kinks or in one of the following three regions, into which the two kinks separate the possible number of positions. First, the optimal n^{*} could be below the first kink at $\left(1 - \lambda\right) n_{-1}$. This means that the firm fires existing workers. Second, n^{*} could be between the two kinks. In this case, the firm fills some but not all of its empty positions, and it does not create any new positions. Third, if n^{*} is above the second kink at n_{-1}, the firm creates new positions and posts vacancies to fill all empty positions.

Let $J\left(n_{-1}, x, s\right) := \frac{\partial V\left(n_{-1}, x, s\right)}{\partial n_{-1}}$ denote the marginal value of an additional position at the beginning of the period. Then, I obtain the following optimal choice of the firm and the corresponding marginal value of a position in each of the five cases:

1. $n^{*} < \left(1 - \lambda\right) n_{-1}$. The firm fires workers. In this case, the number of positions is reduced until marginal profit in the current period and the expected future marginal profit sum to zero, that is $J\left(n_{-1}, x, s\right) = 0$.

---

[12]I use the shorthand notation $\left(y\right)_{+} := \max\left\{y, 0\right\}$.

13

The optimal number of positions n^* is implicitly given by

$$0 = \frac{\partial \Pi}{\partial n} (n^*, x, s) + \beta (1 - \delta) \mathbb{E} \left[J \left((1 - \delta) n^*, x', s' \right) | x, s \right]. \qquad (3)$$

2. $n^* = (1 - \lambda) n_{-1}$. The firm keeps all its existing workers, but does not fill any empty positions. In this case, the marginal value of a worker must be positive but less than the hiring cost $\frac{c}{q}$ so that the firm neither wants to fire workers nor fill empty positions. The marginal value of beginning the period with an additional position created in the past, being filled with $(1 - \lambda)$ workers, is given by the sum of current marginal profit of these workers and the expected future marginal value of additional $(1 - \lambda)(1 - \delta)$ positions in the next period:

$$
\begin{aligned}
J(n_{-1}, x, s) = (1 - \lambda) \frac{\partial \Pi}{\partial n} (n^*, x, s) \\
+ \beta (1 - \lambda)(1 - \delta) \mathbb{E} \left[J \left((1 - \delta) n^*, x', s' \right) | x, s \right]. \quad (4)
\end{aligned}
$$

3. $n^* \in ((1 - \lambda) n_{-1}, n_{-1})$. The firm fills some of its empty positions and shuts down the others. The optimal number of positions is determined by the condition that the marginal current and future expected profit of an additional worker equals its hiring cost:

$$\frac{\partial \Pi}{\partial n} (n^*, x, s) + \beta (1 - \delta) \mathbb{E} \left[J \left((1 - \delta) n^*, x', s' \right) | x, s \right] = \frac{c}{q(\theta(s))}. \qquad (5)$$

The marginal value of an additional position at the beginning of the period is equal to the saved cost of hiring $(1 - \lambda)$ workers:

$$J(n_{-1}, x, s) = (1 - \lambda) \frac{c}{q(\theta(s))}. \qquad (6)$$

4. $n^* = n_{-1}$. The firm keeps and fills all existing empty positions, but it does not create new positions. In this case, the marginal value of a position must exceed the hiring cost for filling but be less than $H + (1 - \lambda) \frac{c}{q}$ so that it lies below the cost of creating and filling positions.

The marginal value of an additional position created in the past is given by the sum of current marginal profit of the marginal worker (including the filling cost for the empty positions) and the expected future marginal value of additional $(1 - \delta)$ positions in the next period:

$$J(n_{-1}, x, s) = \frac{\partial \Pi}{\partial n}(n^*, x, s) - \lambda \frac{c}{q(\theta(s))}$$
$$+ \beta(1 - \delta) \mathbb{E}\left[J\left((1 - \delta)n^*, x', s'\right) | x, s\right] \quad (7)$$

5. $n^* > n_{-1}$. The firm opens new positions and fills all empty positions. Its optimal size is determined by the condition that the cost of creating and filling an additional position equals its marginal current and expected future marginal profit:

$$\frac{\partial \Pi}{\partial n}(n^*, x, s) + \beta(1 - \delta) \mathbb{E}\left[J\left((1 - \delta)n^*, x', s'\right) | x, s\right] = H + \frac{c}{q(\theta(s))}. \quad (8)$$

The marginal value of a position created in the past, filled with $(1 - \lambda)$ workers, is the sum of the saved creation cost and filling cost:

$$J(n_{-1}, x, s) = H + (1 - \lambda) \frac{c}{q(\theta(s))}. \quad (9)$$

In summary, the marginal value function has to satisfy

$$J(n_{-1}, x, s) = \min\left\{H + \frac{(1 - \lambda)c}{q(\theta(s))}, \max\left\{J_F(n_{-1}, x, s),\right.\right.$$
$$\left.\left.\min\left\{\frac{(1 - \lambda)c}{q(\theta(s))}, \max\left\{J_K(n_{-1}, x, s), 0\right\}\right\}\right\}\right\}. \quad (10)$$

where J_K and J_F denote the right hand side values of equations (4) and (7): J_K is the marginal value that results, if the firm keeps all its workers but is inactive with respect to filling positions ($n^* = (1 - \lambda)n_{-1}$), and J_F is the marginal value, if the firm fills all empty positions but is inactive with respect to creating new positions ($n^* = n_{-1}$).

15

The optimal behaviour of a firm can be summarised by three cutoffs, that determine how many positions firms want to create, how many positions they want to fill, and how many workers they want to fire. These cutoffs are functions of idiosyncratic productivity x and the aggregate state s. If the number of positions at the beginning of the period is small, a firm will create positions until the number of positions is equal to the lower bound $n_{LB}(x, s)$. This cutoff is implicitly given by the condition that the marginal value J_F of a firm that fills all position is exactly equal to the upper bound of the marginal value function:

$$J_F(n_{LB}(x, s), x, s) = H + (1 - \lambda) \frac{c}{q(\theta(s))}. \tag{11}$$

If the initial number of positions lies above the lower cutoff, the firm will be inactive with respect to creating new positions but it will fill all empty positions as long as the marginal value J_F is greater than $(1 - \lambda)\frac{c}{q}$. This condition implicitly defines the middle cutoff $n_{MB}(x, s)$:

$$J_F(n_{MB}(x, s), x, s) = (1 - \lambda) \frac{c}{q(\theta(s))}. \tag{12}$$

The third cutoff, $n_{UB}(x, s)$, determines the upper bound of positions in the firm. The firm does not fire existing workers, as long as the marginal value J_K is positive. The upper bound is implicitly defined by

$$J_K(n_{UB}(x, s), x, s) = 0. \tag{13}$$

The cutoffs are increasing in x as higher idiosyncratic productivity makes the marginal position more valuable in a given aggregate state. Figure 3 shows an example of a marginal value function together with the three cutoffs and the optimal firm decision. Note that for small λ in a weekly calibration, the third case described above almost collapses and the firm usually fills either all or none of its empty positions.
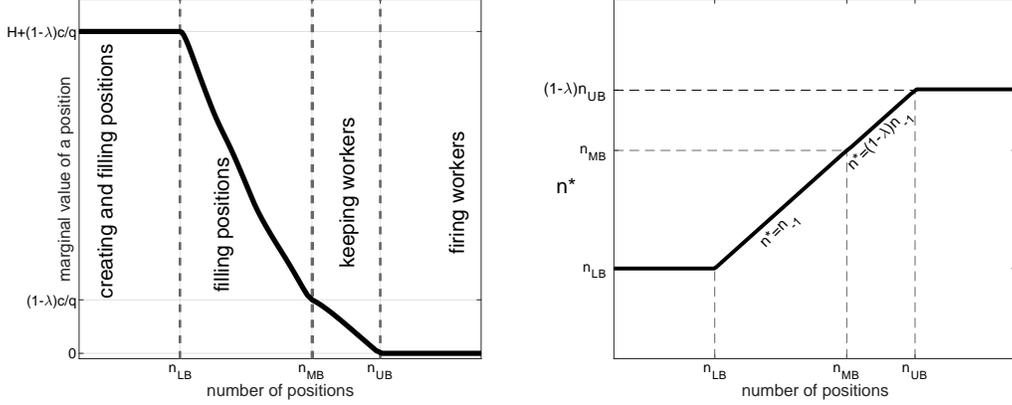
Figure 3: Marginal value function and optimal policy
Notes. The left panel shows the marginal value of a position at (including $1 - \lambda$ workers) for a given level of idiosyncratic productivity x and aggregate state s as a function of the number of positions at the firm. The right panel shows the corresponding optimal number of filled positions in the production phase.

3.3 Market clearing

In equilibrium, the following market clearing condition has to hold: the mass of workers that firms want to hire must be equal to the number of newly matched unemployed according to the matching function. Let $F_s(n_{i,-1}, x_i)$ denote the distribution of positions and idiosyncratic productivity across firms in the aggregate state s. Then

$$\int_{n_{i,-1}} \int_{x_i} \left(n^*(n_{i,-1}, x_i, s) - (1 - \lambda) n_{i,-1}\right)_+ dF_s(n_{i,-1}, x_i)$$

$$= M(u(s), v(s)) \quad (14)$$

must hold in all possible states s, where $M(u, v)$ is the number of matches when the mass of unemployed is u and the mass of posted vacancies is v. It is important to ensure that this condition is satisfied when the economy is simulated in the numerical solution. In Appendix F, I describe the algorithm that I develop to ensure market clearing exctly.[13]

---

[13]This is important because differences between aggregate supply and aggregate demand might accumulate over time in a simulation.

4 Parametrization

The model is parametrized to US data. In section 6, I compare the results from the model with the corresponding time series in the US. Because the probability for unemployed to find a job within a month is close to 50%, I choose a high frequency for my model. This allows unemployed to find a job in less than a month. In particular, I use a weekly parametrization like Elsby and Michaels (2013) with a discount factor $\beta = 0.999$.[14]

I calibrate the model for various values of the cost of creating positions H. In the next sections, I will compare the results for those different values of H. Most of the parameters are held constant across calibrations, and they are summarised in Table 1. There are three parameters that I recalibrate in order to meet the following three targets across calibrations: in the absence of aggregate shocks, the unemployment rate should be 6.5%, the job-finding rate should be 11.25%, and the cost of filling a vacancy should be 14% of the average quarterly wage.[15] In order to meet these targets, I choose ω_3 in the wage equation, the standard deviation $\bar{\sigma}_x$ of the idiosyncratic shock, and the mass of firms, as described below. The resulting parameters for different values of H are listed in Table 2.

Matching function and vacancy costs. The matching function is assumed to take the Cobb-Douglas form $M(u,v) = \mu u^\eta v^{1-\eta}$. The elasticity η is set to 0.6 as in Elsby and Michaels (2013), based on estimates in Petrongolo and Pissarides (2001). This implies that the job-finding probability for an unemployed as a function of market tightness is $a(\theta) = \mu\theta^{1-\eta}$. The probability of a vacancy being filled is $q(\theta) = \mu\theta^{-\eta}$. Without loss of generality, I target market tightness to be equal to 0.72 in the absence of aggregate shocks.[16] The scale of the matching function $\mu = 0.128$ is implied by the targeted job-finding

---

[14]This is equivalent to an annual real interest rate of 5.3%.

[15]These targets are also used by Elsby and Michaels (2013).

[16]Note that the scale of the matching function μ and the vacancy posting cost c can be adjusted in response to a proportional change in market tightness such that $a(\theta)$ and $c\theta$ remain unchanged. Then, the job-finding rate remains the same and the firm's problem is not affected because the filling cost $\frac{c}{q(\theta)} = \frac{c\theta}{a(\theta)}$ and the term $c\theta$ in the wage equation are unchanged.

Table 1: Model parameters

Parameter	Value	Description
β	0.999	Discount factor
η	0.6	Elasticity of matching function ($M(u,v) = \mu u^\eta v^{1-\eta}$)
μ	0.128	Scale of matching function
c	0.156	Vacancy posting cost
ω_1	0.8	Firm's share of production
		(Stole-Zwiebel bargaining with bargaining power 0.5)
ω_2	$0.5c$	Part of the wage that is proportional to θ
δ	0.001	Depreciation of positions
λ	0.003	Exogenous worker turnover
α	0.75	Elasticity of production function
H	$0-10$	Cost of creating a position (relative to filling it)
λ_x	0.043	Arrival rate of idiosyncratic shock
σ_σ	0.08	Unconditional standard deviation of $\log \bar{\sigma}_x$
ρ_σ	0.988	Autocorrelation of $\log \bar{\sigma}_x$
σ_z	0.02	Unconditional standard deviation of $\log z$
ρ_z	0.996	Autocorrelation of $\log z$

Notes. Parameters are based on a weekly calibration. The model is solved for different values of creation cost H. The parameters that are recalibrated depending on H can be found in Table 2.

probability $a(0.72) = 0.1125$. Vacancy posting cost c is chosen to normalise the cost of filling positions to 1 ($c = q(0.72)$). Therefore, the creation cost H can be interpreted as a multiple of the filling cost in steady state.

Production function. The production function takes the form $F(n) = n^\alpha$. I set $\alpha = 0.75$ in between the values used by Elsby and Michaels (2013) and Schaal (2017). Creation costs could be interpreted as the value of capital that is lost if a position is shut down and the respective share of capital is sold. Therefore, α should not be interpreted as just the curvature of the production function when increasing the number of workers, holding capital constant. More precisely, $F(n)$ is the revenue function rather than the amount of goods produced by the firm. A downward sloping demand curve, for example due to monopolistic competition, reduces α relative to the curvature in the pure production function.

Table 2: Calibration results

H	0	5	10	Cost of creating a position
Calibrated parameters				
ω_3	0.354	0.346	0.338	Constant in wage equation
$\bar{\sigma}_x$	0.131	0.351	0.502	Standard deviation of idiosyncratic shock
$mass$	0.197	0.156	0.116	Mass of firms
Non-targeted statistics without aggregate risk				
	0%	52%	99%	Creation cost / quarterly output
	0%	4.1%	7.4%	Aggregate creation costs / output
	1.2%	1.1%	1.1%	Aggregate filling costs / output
	0.79	0.74	0.71	Labour share
	0.40	0.41	0.37	Std. dev. of annual employment growth

Notes. List of parameters that are recalibrated depending on H to keep the unemployment rate, job-finding rate, and filling cost relative to average wage the same across calibrations. The lower part reports some statistics of the model without aggregate risk. The last row reports the cross-sectional standard deviation of annual employment growth calculated as $std\left(\frac{n_{i,t}-n_{i,t-52}}{0.5(n_{i,t}+n_{i,t-52})}\right)$.

Wage equation. I impose the following functional form for the total wage bill:[17]

$$W\left(n,x,s\right)=\left(1-\omega_1\right)zxn^{\alpha}+\left(\omega_2\theta+\omega_3\right)n. \tag{15}$$

This form is common in search and matching models:[18] workers get a share

---

[17]The total wage bill is specified as opposed to the wage per worker to simplify notation. The derivative of W with respect to n gives the effect of hiring an additional worker on a firm's wage bill. This differs from the wage per worker because the wage per worker is decreasing in firm size. Therefore, hiring an additional worker reduces the wage a firm has to pay to its existing workers, dampening the increase of the total wage bill.

[18]It can be derived by imposing Stole-Zwiebel bargaining in the absence of creation costs for positions ($H=0$) like in Elsby and Michaels (2013). There is no analogous analytical solution to the bargaining problem in my full model when $H>0$. This is because some workers are hired at firms that newly create positions, and others are hired to fill positions that have become empty due to exogenous worker turn-over. The surplus of a match is not equal across these different positions. Therefore, one needs to know the distribution of vacancies across firms, if one wants to calculate the expected value of an unemployed, for instance. As a result, there is no analytical solution to the bargaining problem. To make the bargaining problem simpler, it could be assumed that firms have all the bargaining power. Then, the value for a newly employed worker does not exceed the value of an unemployed. This approach was used for example in Cooper et al. (2007) and in Postel-Vinay and Robin

of the marginal product and a payment that compensates them according to their outside option. This depends on the value of home production and on the option value of finding another job, which increases in market tightness.

The parameters ω_1 and ω_2 are set to the respective values derived from Stole and Zwiebel (1996) bargaining when $H = 0$ and the worker's bargaining power is 0.5. The resulting values are $\omega_1 = 0.8$ and $\omega_2 = 0.5c$.[19] The constant ω_3 is calibrated depending on H such that the calibration targets are met.

Attrition. I choose the depreciation rate of positions δ to be 0.001 on a weekly basis, which corresponds to 5% on a yearly basis. The choice of λ determines the exogenous rate of workers leaving a firm. It should represent the rate at which employees quit for reasons other than low productivity of the firm. Thus it also determines the share of vacancies that are posted by firms to replace these workers relative to vacancies posted to fill newly created positions. Given that the total separation rate is implied by the calibrated values for unemployment and the job-finding rate, λ also determines how much job destruction is exogenous and endogenous, respectively. The probability of a worker leaving for exogenous reasons is approximately $\lambda + \delta$. The average separation rate in the model is 0.008. I set $\lambda = 0.003$ such that about half of separations are exogenous in the absence of aggregate shocks. This corresponds to quits being approximately half of total separations in the US.[20]

_____

(2002). In my framework, it would imply that all workers are paid the same wage equal to the flow benefits of an unemployed.

[19]It is not crucial that the wage per worker differs across firms, i.e. that ω_1 is less than 1. A firm's production less wages is equal to $\omega_1 z x n^\alpha - (\omega_2 \theta + \omega_3) n$. From a firm's perspective, changing only ω_1 to be 1 is equivalent to a permanent change in productivity. The firm optimally responds by multiplying its number of positions by the factor $\omega_1^{-\frac{1}{1-\alpha}}$. Then all first order conditions hold exactly as before. This means that all firms become proportionally larger which can be perfectly offset by multiplying the mass of firms by $\omega_1^{\frac{1}{1-\alpha}}$. Then the employment dynamics of the model are identical regardless of the value ω_1 even with aggregate risk. The average wage, however, is affected by a different choice of ω_1. As cost of filling relative to the average wage is targeted in the calibration, the parameter ω_3 has to be recalibrated when ω_1 is changed. Therefore, the results of the calibrated model do depend on the choice for ω_1. Figure B.6 in Appendix B.1 shows that the main results are robust to choosing $\omega_1 = 1$.

[20]Monthly quits are on average 52% of total separations from 2001-2013 in the JOLTS dataset produced by the BLS. Note that while the probability of an exogenous quit in my model is time-invariant, the share of quits in total separations is procyclical because the

21

Idiosyncratic shocks. Like in Elsby and Michaels (2013), idiosyncratic shocks arrive at rate $\lambda_x = 0.043$. When a firm is hit by a shock, the newly drawn idiosyncratic productivity follows a log-normal distribution:

$$\log x_t^{new} \sim N\left(0, \sigma_{x,t}^2\right). \tag{16}$$

Standard deviation of idiosyncratic productivity in steady state ($\bar{\sigma}_x$) is recalibrated when H is varied. This is necessary because a higher cost of creation makes firms more reluctant to destroy positions and fire workers. Hence, endogenous job destruction falls. Table 2 reports the resulting standard deviation that is increasing in H. Note that this does not necessarily mean that the number of positions within a firm becomes more volatile. The last row of Table 2 shows that for all values of H, the standard deviation of annual employment growth is close to the estimate of 0.416 calculated by Elsby and Michaels (2013) using data on continuing establishments in the Longitudinal Business Database from 1992 to 2005.[21]

Uncertainty shocks. The standard deviation of idiosyncratic productivity follows a log-normal process:

$$\log \sigma_{x,t} - \log \bar{\sigma}_x = \rho_\sigma \left(\log \sigma_{x,t-1} - \log \bar{\sigma}_x\right) + \varepsilon_{\sigma,t}. \tag{17}$$

When idiosyncratic volatility changes, I also adjust aggregate productivity to compensate for the following effects. As idiosyncratic productivity follows a log-normal distribution, increases of $\sigma_{x,t}$ increase average productivity across firms.[22] In addition, if firms wanted to keep their marginal product the same, the specified production function implies that labour demand is a convex function of idiosyncratic productivity x.[23] Therefore, higher dispersion of productivity across firms increases aggregate demand for labour. I abstract from this

---

separation rate is countercyclical. It would be a simple extension of the model to make exogenous separations dependent on the aggregate productivity and uncertainty states.

[21]The annual growth rate is calculated as $\frac{n_{i,t} - n_{i,t-52}}{0.5(n_{i,t} + n_{i,t-52})}$ as in Davis and Haltiwanger (1992).

[22]The expected newly drawn level of productivity is given by $\mathbb{E} x_t^{new} = \exp\left(\frac{\sigma_{x,t}^2}{2}\right)$.

[23]This can be seen in a simple static model, in which the marginal product of a worker

effect by multiplying productivity with $\tilde{z}_t$, dependent on the cross-sectional distribution of idiosyncratic productivity, such that labour demand would not be affected by changes in uncertainty in a "perfect insurance" model:[24]

$$\int_{n_{i,-1}} \int_{x_i} (\tilde{z}_t x_{i,t})^{\frac{1}{1-\alpha}} \, dF_s\left(n_{i,-1}, x_i\right) = const. \tag{18}$$

Aggregate productivity shocks. Aggregate productivity z follows a log-normal process that is independent from the volatility of idiosyncratic productivity:

$$\log z_t = \rho_z \log z_{t-1} + \varepsilon_{z,t}. \tag{19}$$

The weekly autocorrelation ρ_z is chosen to be 0.996, which corresponds to a quarterly autocorrelation of 0.95. The standard deviation of the innovation is chosen such that the unconditional standard deviation of log-productivity is 0.02.

Creation cost. I solve the model for various values of H ranging from 0 to 10. Table 2 shows that this means that the cost of creating a position is about $10 \cdot H\%$ of quarterly output. Thus, for the highest cost considered ($H = 10$) it is about as costly for firms to create one position as the average quarterly output of each position is. I also report the share of output that firms spend on posting vacancies and on creating positions in Table 2. This filling cost is a bit more than 1% of output, and the aggregate creation cost ranges from 0 to 7.4% of output. Adding the creation costs, filling costs and the labour share in Table 2, it can be seen that average firm profits are about 20% of output. They

---

equals the wage:

$$\alpha z x_i n_i^{\alpha-1} = w.$$

Then, labour demand by firm i is given as

$$n_i = \left(\frac{\alpha z x_i}{w}\right)^{\frac{1}{1-\alpha}}$$

[24] "Perfect insurance" means that firms are always at the creation cutoff $n_{LB}(x)$. This would happen if it was possible to trade existing positions and workers between firms, until the marginal product at all firms equalises. For the numerical solution of the model, this also means that $m^x \equiv \int_i x_i^{\frac{1}{1-\alpha}} \, dF(i)$ becomes a state variable.

could be interpreted as the return on capital, which is absent in my model. Then, the implied capital per position is about 15 times quarterly output.[25] Hence, the largest creation costs $H = 10$, considered in this paper, could be interpreted in the following way: when a position is created the firm has to pay about 7% of the value of capital that is needed for the position. This amount cannot be recovered when the position is destroyed and the capital is sold.[26]

5 Effects of uncertainty in the model

In the first subsection, I study the long run effects of a permanent change in idiosyncratic volatility. It shows that the creation cost H is crucial for the response of the job-finding rate after an increase in uncertainty. When $H = 0$, the job-finding rate is hardly affected, but when $H = 10$, it falls significantly and this is responsible for almost half of the drop in employment. In the second subsection, I decompose the effects of higher uncertainty into changes driven by the wait-and-see effect, the realized volatility effect, and the general equilibrium effect as market tightness changes. It demonstrates that the wait-and-see effect becomes stronger when H is large. This is important, because the wait-and-see effect ultimately reduces the job-finding rate. In subsection 5.3, I solve the full model with time-varying aggregate productivity and time-varying uncertainty. The impulse response functions after an uncertainty shock confirm that it is only in case of high creation cost that the job-finding rate falls significantly. I also show that the quantitative effect of an increase in

---

[25]This is calculated using the quarterly discount factor, which implies the quarterly interest rate $\frac{\beta^{13}}{1-\beta^{13}}$, and a profit share of 20%. Then the implied ratio of capital and quarterly output is $0.2\frac{\beta^{13}}{1-\beta^{13}}$.

[26]The resale price of capital in case of partial irreversibility is usually calibrated to match some business cycle statistics. There is a wide range of values obtained. For example, Khan and Thomas (2013) use an investment resale loss of 4.6%, whereas Bloom (2009) has a resale loss of 33.9% in his full model. An alternative way to compare the creation cost used in this paper with the capital adjustment cost literature is to calculate the aggregate amount of capital adjustment costs spent in the economy in equilibrium. Cooper et al. (2004) estimate capital disruption costs that are 15% of revenues minus labour costs. An analogous estimate of the aggregate creation costs in my model is 16% when $H = 5$.

uncertainty depends on the state of the economy. In particular, higher uncertainty increases unemployment by more when the initial unemployment rate was already high.
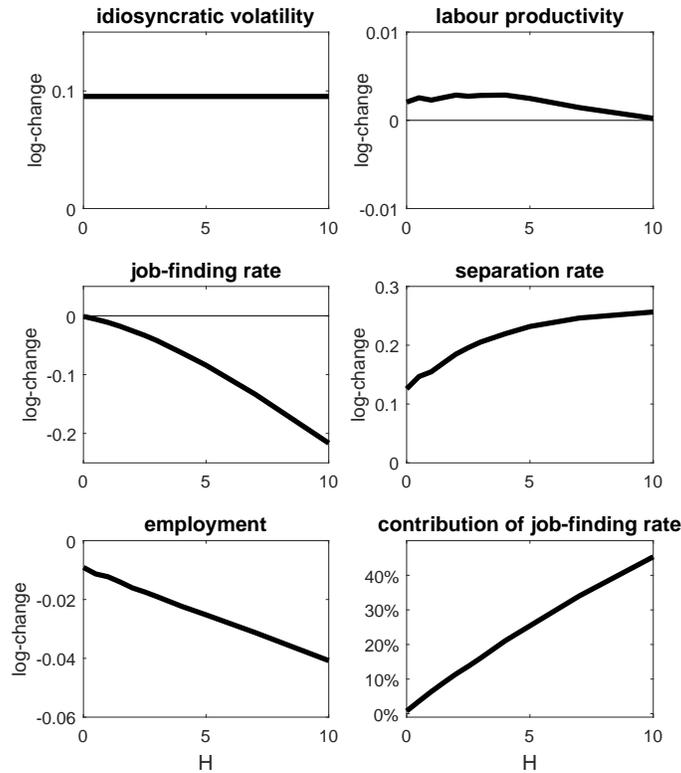
5.1 Comparative statics

In this section, I consider the long run effects of a permanent increase in idiosyncratic volatility in steady state. I demonstrate that the presence of creation costs is necessary to get a fall in the job-finding rate as found in the empirical part. Figure 4 shows the effect of a 10% increase in idiosyncratic volatility σ_x for increasing values of cost H.[27] This second moment shock is accompanied with a negative first moment shock, as described above, such that labour productivity would not be affected in a "perfect insurance" environment, in which the marginal product is equal at all firms.

From the third subplot it becomes evident that without creation costs ($H = 0$), the job-finding rate hardly changes in response to higher uncertainty. When H becomes larger, the job-finding rate falls significantly. The separation rate also reacts more strongly with higher creation costs. Both, the higher separation rate and the lower job-finding rate imply a stronger fall in employment after a permanent increase in idiosyncratic volatility for high H. In the sixth subplot I decompose how much of the change in unemployment is due to the higher separation rate and how much is due to the lower job-finding rate.[28]

---

[27]Note that due to the calibration strategy, idiosyncratic volatility is increasing in H to get the same level of endogenous job destruction. Therefore, a 10% increase, raises idiosyncratic volatility by more percentage points with higher creation costs, which is a reason for the stronger response of the unemployment rate. My focus in interpreting the results, however, is on the relative contribution of job-finding rate and separation rate. Figure B.2 in Appendix B.1 shows the effects, when idiosyncratic volatility is increased by the same absolute amount. Then, the response of employment is comparable across calibrations. Importantly, the contribution of the job-finding rate to the fall in employment is hardly affected by the alternative size of the shock.

[28]I calculate the change in unemployment if only the job-finding rate was changed, Δu^{JFR}, and the change in unemployment if only the separation rate was changed, Δu^{SR}. Note that $\Delta u \approx \Delta u^{JFR} + \Delta u^{SR}$ does not hold exactly. I report the contribution of the job-finding rate as $\frac{\Delta u^{JFR}}{\Delta u^{JFR} + \Delta u^{SR}}$.

Figure 4: Comparative statics

Notes. Responses after a permanent 10% increase of idiosyncratic volatility for varying H. The job-finding rate only falls and thus contributes to lower employment in the presence of creation cost $H > 0$.

There are two opposing effects that change labour productivity in this heterogeneous agents model. First, the wait-and-see effect increases dispersion of firms conditional on idiosyncratic productivity, because they become more reluctant to create and to destroy positions. This makes the average product of labour more dispersed across firms and leads to a reduction of measured labour productivity in the aggregate.[29] The second effect is that reduced average employment increases the average product of firms because they operate at decreasing returns to scale. This effect increases labour productivity. In the

---

[29]This effect on measured TFP has also been discussed by Bloom et al. (2018).

comparative statics exercise the latter effect dominates. In the full model, however, labour productivity can fall in response to an increase in uncertainty when H is large, as can be seen in Figure 7.

5.2 Decomposition of effects

This section decomposes the effects of higher uncertainty into three components, namely the wait-and-see effect, the realized volatility effect and the general equilibrium effect due to changes in market tightness.

The wait-and-see effect means that higher uncertainty makes firms more reluctant to create positions, because a future adverse productivity shock becomes more likely. This reduces the expected continuation value and and shifts the creation cutoff n_{LB} down. Likewise, firms fire fewer workers as the chance of a future positive productivity shock becomes higher. The resulting increase of the expected continuation value shifts the firing cutoff n_{UB} up. Hence, the wait-and-see effect implies that higher uncertainty leads to fewer positions being created and fewer workers being fired.

The realized volatility effect occurs, because a higher volatility of idiosyncratic firm productivity makes more firms experience large positive or negative shocks. Holding a firm's cutoffs constant, it is thus more likely to get a low enough productivity draw that makes it fire workers, and also more likely to get a high enough productivity draw that makes it create new positions. Therefore, the realized volatility effect leads to more positions being created and more workers being fired in the presence of higher uncertainty.

Thirdly, in general equilibrium changes in market tightness affect a firm's optimal behaviour. For example, if the combination of wait-and-see effect and realized volatility effect imply that the aggregate demand for new hires falls short of supply, market tightness falls to ensure market clearing in the matching market. Then, it becomes cheaper for firms to fill their empty positions. This primarily shifts the intermediate cutoff n_{MB} upwards, which means that they are more willing to fill an empty position as opposed to shutting it down. In the following subsections, I discuss the three components in detail.

5.2.1 Wait-and-see effect

The wait-and-see effect captures how much economic variables are affected by a change in the optimal cutoffs of firms in response to higher idiosyncratic uncertainty. I construct a measure for how much the cutoffs differ from their "perfect-insurance" counterparts. Under "perfect-insurance" firms would not be affected by idiosyncratic shocks, because they would be allowed to sell their existing positions and workers to other firms without any adjustment cost. The deviation of cutoffs from this solution measures how much idiosyncratic risk affects a firm's decisions. The method is explained in detail in Appendix C.1.
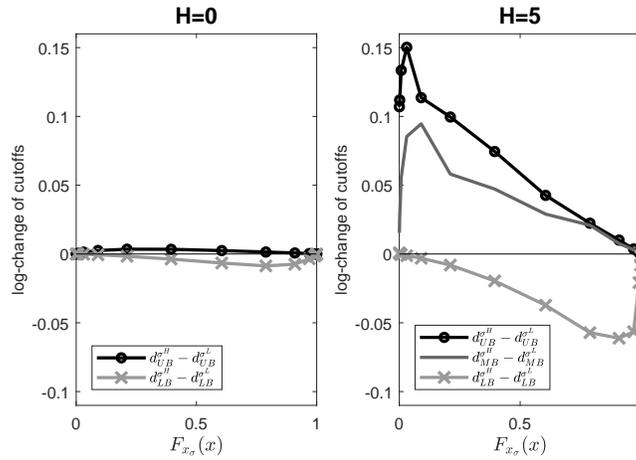
Figure 5: Wait-and-see effect

Notes. The figure draws the change in cutoffs after a 10% increase in uncertainty relative to their "perfect-insurance" counterparts. These changes are drawn as a function of the percentile in the idiosyncratic productivity distribution ($F_{x_\sigma}(x)$). Note that the lower cutoff, determining the creation of positions, and the middle cutoff, determining whether empty positions are filled, coincide in the absence of creation cost ($H = 0$).

Figure 5 shows how these three cutoff functions change, when idiosyncratic volatility is increased. The cutoffs at the lower bound fall, which means that comparable firms will create fewer positions in the presence of higher uncertainty. In contrast to that, the cutoffs at the upper bound increase, which makes firms fire fewer workers.[30]

---

[30] Appendix C.1 discusses the reasons for the hump-shaped shift of cutoffs in Figure 5.

The reason is that the marginal continuation value is concave with respect to idiosyncratic productivity around the creation cutoff: if a firm creates positions in the current period, the marginal value of a position is at the upper bound. Higher idiosyncratic productivity would therefore not increase the marginal value while lower productivity would reduce it. Therefore, an increase in the second moment of idiosyncratic productivity reduces the expected continuation value of a position and reduces the number of created positions.

The analogous reasoning applies at the upper bound. There, the continuation value is locally convex as the marginal value is at its lower bound when a firm fires workers in the current period. Then, an increase of the second moment increases the expected continuation value and reduces the number of fired workers.

The direction of the shift at the medium cutoff is ambiguous. As the medium cutoff is determined by the condition whether a firm wants to fill empty positions or not, there are two opposing effects. On the one hand, the firm incurs a cost if it hires new workers. This is similar to a firm at the lower cutoff which creates positions and hires new workers. Higher uncertainty makes firms more reluctant to hire new workers, which would shift down the cutoff. On the other hand, if the firm decides not to hire new workers, it shuts down the empty positions. This is similar to a firm at the upper cutoff, which not only fires workers but also shuts down empty positions. Higher uncertainty makes firms also more reluctant to do so. This would shift the medium cutoff upwards. Whether the medium cutoff shifts upwards or downwards is then determined by the size of the creation cost relative to the filling cost. The higher the cost to create a position, the more reluctant a firm is to shut it down. Therefore, for sufficiently high H, the cutoff will shift upwards as can be seen for $H = 5$ in Figure 5. In contrast, when there is only a small cost to create positions, the reluctance of the firm to fill empty positions becomes dominant, and the cutoff will shift down. In particular, in the absence of creation cost ($H = 0$), the medium and lower cutoffs coincide such that they move down as drawn in the left panel of Figure 5.

Using the change of cutoffs, it is possible to calculate the implications of the wait-and-see effect isolated from the other effects when uncertainty increases. Shifting each of the three cutoffs by the respective change calculated above, I obtain the cutoffs $n_j^{W\&S}$ for the wait-and-see effect:

$$\log n_j^{W\&S}\left(F_{\sigma_x}^{-1}(p)\right) = \log n_j\left(F_{\sigma_x}^{-1}(p)\right) + \left[d_j^{\sigma^H}(p) - d_j^{\sigma^L}(p)\right],$$

$$j \in \{LB, MB, UB\}. \quad (20)$$

This means that the cutoffs are shifted by the amount shown in Figure 5, but there is no actual change in volatility when the economy is simulated in partial equilibrium. The latter effect is the realized volatility effect described in the next subsection.

The light grey, dashed lines in Figure 6 show the change of employment, separation rate and demand and supply of newly hired workers after a 10% increase in uncertainty as a function of creation cost H. The separation rate falls unambiguously, because the firing cutoff shifts upwards. The size of the effect is increasing in H, because the wait-and-see effect becomes stronger as argued above. The aggregate demand for newly hired workers falls as well. This is driven by the downward shift of the creation cutoff $n_{LB}(x)$.[31]

Taking the changes of the separation rate and of the number of new hires together, one can calculate the effect on employment. As the two effects run in opposite directions, the resulting change in employment is ambiguous. It turns out that the reduced separation rate dominates for high H and employment rises slightly. In the fourth subplot of Figure 6, the implied change of the supply of newly hired is drawn. Note that market tightness remains unchanged in partial equilibrium such that changes in supply are entirely driven by an increase or decrease of the pool of unemployed. Therefore, whenever employment increases, there are fewer unemployed and the matching function implies that the number of new matches changes proportionally to this change

---

[31]Note that the middle cutoff $n_{MB}(x)$ also determines how many firms want to fill their empty positions. This cutoff shifts up for larger H as seen above. The willingness of more firms to fill their vacant positions increases demand for new hires. But it is dominated by the change in the creation cutoff.
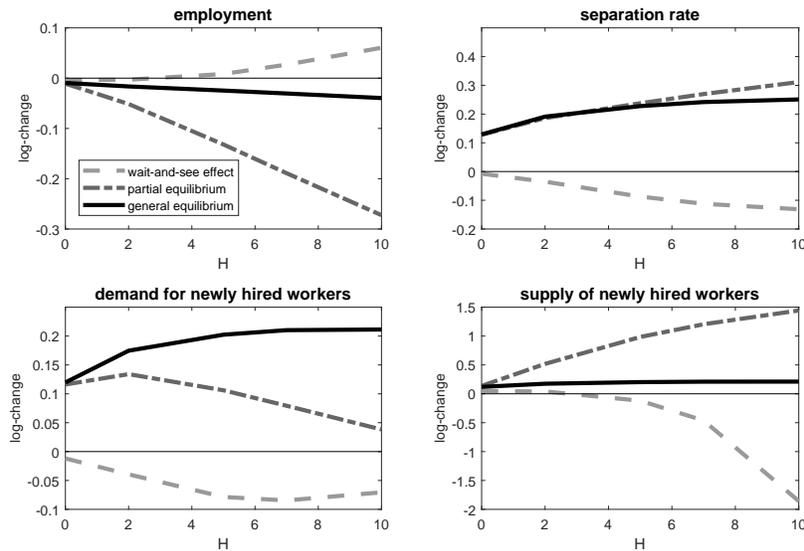
Figure 6: Decomposition exercise

Notes. Decomposition of a 10% increase in uncertainty into wait-and-see effect (light gray curves), realized volatility effect (difference between dark gray and light gray curves), and general equilibrium effect (difference between black and dark grey curves).

in unemployment.

5.2.2 Realized volatility effect

The realized volatility effect captures the effects of higher uncertainty through the increase in the cross-sectional dispersion of idiosyncratic productivity. When volatility increases, there are more firms that get hit by sufficiently negative or positive shocks that lead them to fire workers or to create new positions. Hence, the realized volatility effect captures the changes due to larger shocks occuring, holding the cutoffs fixed, whereas the wait-and-see effect accounts for the change of cutoffs. The combined impact of wait-and-see effect and realized volatility effect is given by the partial equilibrium response after an increase in uncertainty. This means that market tightness is held constant, but firms optimally adjust their cutoffs and experience larger shocks. These partial equilibrium responses are drawn as grey dash-dotted lines in Figure 6. The difference between the partial equilibrium effect and the wait-and-see ef-

fect is then due to the realized volatility effect.

If firms get hit by larger shocks, but did not adjust their cutoffs, they are more likely to be below the creation cutoff or above the firing cutoff at the beginning of the period. This can be seen in Figure 6 as an increase in the separation rate and an increase in the demand for newly hired workers. The higher separation rate has a negative effect on employment, whereas more newly hired workers would increase employment. It can be seen from the first subplot that the former effect dominates and employment falls for all values of creation cost. This fall is increasing in H. Lower employment increases the pool of unemployed, which in turn increases the supply of new matches, as can be seen in the fourth subplot. Note that for larger values of H, the increase in the supply of new workers exceeds the increase in demand for new workers. Therefore, market tightness has to fall to get market clearing, as explained in the next subsection.

5.2.3 General equilibrium effect

Market tightness has to adjust to clear the market, when supply of new matches does not change by the same proportion as aggregate demand of firms for new matches after an increase in uncertainty. In the subsection above, it could be seen that higher uncertainty leads to an excess supply of workers. Therefore, market tightness falls in general equilibrium. This reduces supply of new matches, because the job-finding rate for unemployed falls. Aggregate demand increases, because it becomes easier and thus cheaper for firms to fill their empty positions. In addition, lower market tightness reduces the wage and thus increases profits. Figure C.2 in appendix C.3 shows the shift of cutoffs that is due to the change in market tightness.

Figure 6 draws the general equilibrium effects of higher uncertainty. Lower market tightness primarily increases demand for newly hired workers, while the separation rate is hardly affected. The result is that employment increases relative to partial equilibrium. Both lower market tightness and lower unemployment reduce the supply of new matches. In general equilibrium, the change of demand and supply of new hires must be the same. Note that the

excess supply of workers was increasing in H in partial equilibrium. This leads to a bigger fall of market tightness when H is large. As a result, the job-finding rate falls in the presence of a large creation cost as could be seen in Figure 4.[32]

5.3 IRFs of uncertainty shocks

In this subsection, I compute impulse response functions after an increase in idiosyncratic uncertainty in the full model. I show that higher uncertainty leads to a fall in employment. In the absence of creation costs, this fall in employment is entirely driven by a higher separation rate. The higher the creation cost H becomes, the more the job-finding rate falls similar to the results in the comparative statics exercise above. This means that the job-finding rate contributes more to changes in employment.

In the full model with time-varying aggregate shocks, the response of endogenous variables after a shock depends on the distribution of firms when the shock hits. For example, in some periods an increase in uncertainty leads to a fall in employment that is twice as strong as the fall after a shock in other periods. In particular, I show that the response of employment becomes stronger, when a shock hits in periods of high unemployment.

Figure 7 shows the median response of labour market variables for different values of the creation cost.[33] Similar to the results in the comparative statics

---

[32]It is important to distinguish between the job-finding rate and the number of newly hired workers, which is given by the product of job-finding rate and the number of unemployed. While higher uncertainty reduces the job-finding rate when H is large, the absolute number of new matches increases, because unemployment increases. In steady state, unemployment rate u, job-finding rate f, and separation rate s satisfy $u = \frac{s}{s+f}$. Then the change of newly hired workers fu is given by

$$d \log (fu) = (1 - u) \, d \log s + u \, d \log f.$$

As $(1 - u)$ is considerably larger than u, the number of newly hired will increase unless the job-finding rate falls by more than a multiple of the increase in the separation rate:

$$d \log (fu) > 0, \text{ if } -d \log f < \frac{1-u}{u} d \log s.$$

[33]When comparing the IRFs with the ones estimated in the US data in Figure 2, note that a change in idiosyncratic volatility has bigger effects on the transition rates than a change in
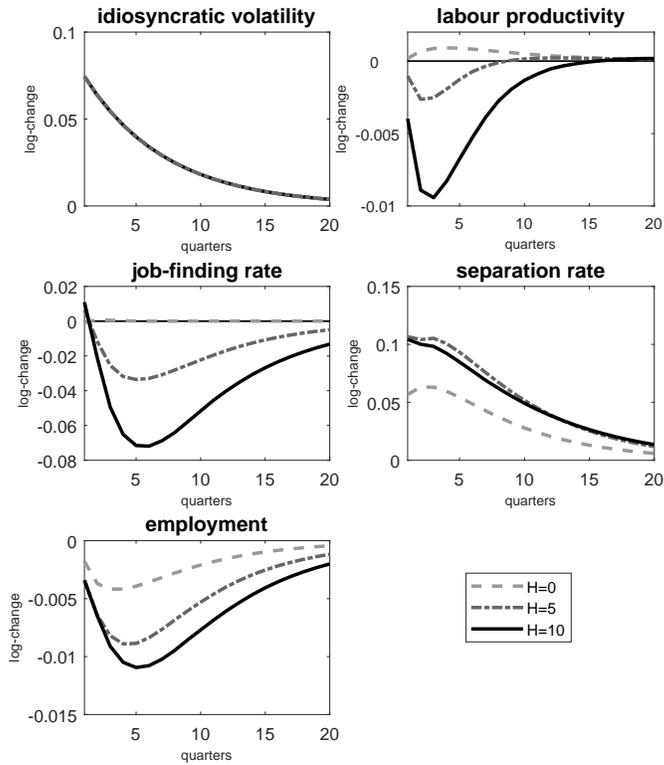
Figure 7: Impulse response functions
Notes. The responses after a one standard deviation increase of σ_x are calculated at 500 different times along a randomly drawn time path of aggregate shocks. The figure draws the median value of these realizations of the respective variable.

exercise above, the job-finding rate hardly responds in the absence of creation costs. In the presence of higher creation costs the job-finding rate falls. The increase in the separation rate becomes stronger as well, but the peak change of the job-finding rate becomes larger relative to the peak change of the separation rate. When $H = 10$, the peak change of the job-finding rate is more

the empirical volatility measure in the data. It is not surprising that the empirical measure obtained from implied volatilities is not directly comparable to changes in the standard deviation of the idiosyncratic shock, which is not observed in the data. Therefore, my focus is on analysing to what extent the separation rate and the job-finding rate drive changes in employment. In section 6, I use a scaled down version of the empirical volatility measure into my model to compare the model's performance with time series data from the US.

than 70% of the peak change in the separation rate. Compared to the US data in Figure 2, the peak fall in the job-finding rate is still too small, as it was found to be about twice as big relative to the increase in the separation rate. The model can capture well the timing of the peak responses. The job-finding rate peaks after around 5 quarters in the model and in the data, whereas the separation rate peaks early after the shock.

Labour productivity increases slightly for low values of the creation cost, while it falls for higher values. This fall is qualitatively different compared to the comparative statics exercise, where labour productivity increased in Figure 4. It changes in the model, because the distribution of positions across firms changes. When the creation cost is higher, the wait-and-see effect becomes stronger. This increases the dispersion of the distribution of positions over firms with the same idiosyncratic productivity, and reduces measured labour productivity. This effect dominates the implications of decreasing returns to scale, which increase the average product when employment falls. Quantitatively, the initial drop in labour productivity for high H is comparable to the IRF estimated in the US data, when the size of the shock is chosen to get the same peak response of employment.

Whereas Figure 7 plots the median impulse response functions in the model, the observed responses can vary significantly depending on the initial state. In the model, the cross-sectional distribution of positions and idiosyncratic productivity is a state variable. As it is solved using non-linear methods, the response to shocks can differ depending on the initial distribution. In particular, for higher values of the creation cost H, the dependence on the initial state becomes more relevant. Figure 8 repeats the median IRFs from Figure 7 for $H = 0$ and $H = 5$. In addition, it also shows the 5^{th} and the 95^{th} percentile of the responses in dashed lines.

In theory, the whole distribution at the initial state matters. One moment of this distribution, which is easily observable, is the unemployment rate when the shock hits. On average, employment falls more after an increase in uncertainty when the unemployment rate was already high at the time the shock hits as is shown in Appendix D.
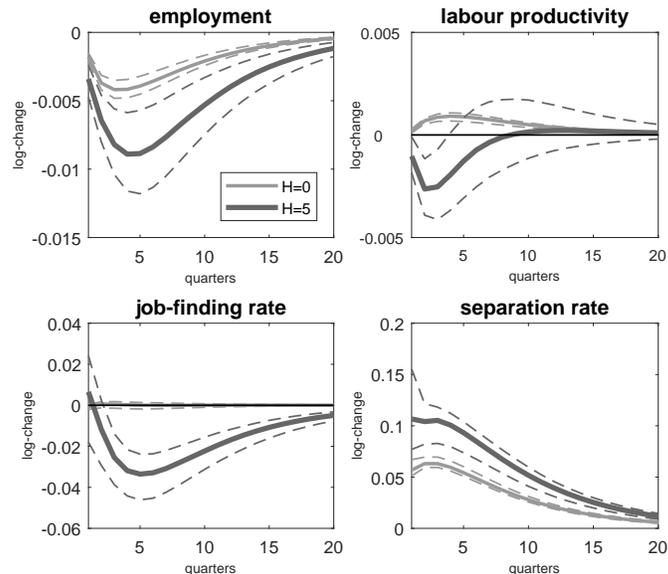
Figure 8: State dependence of impulse response functions
Notes. Impulse response functions for a one standard deviation idiosyncratic volatility shock in the model. Solid lines depict the median change of the respective variable. Dashed lines draw the 5^{th} and the 95^{th} percentiles of the responses.
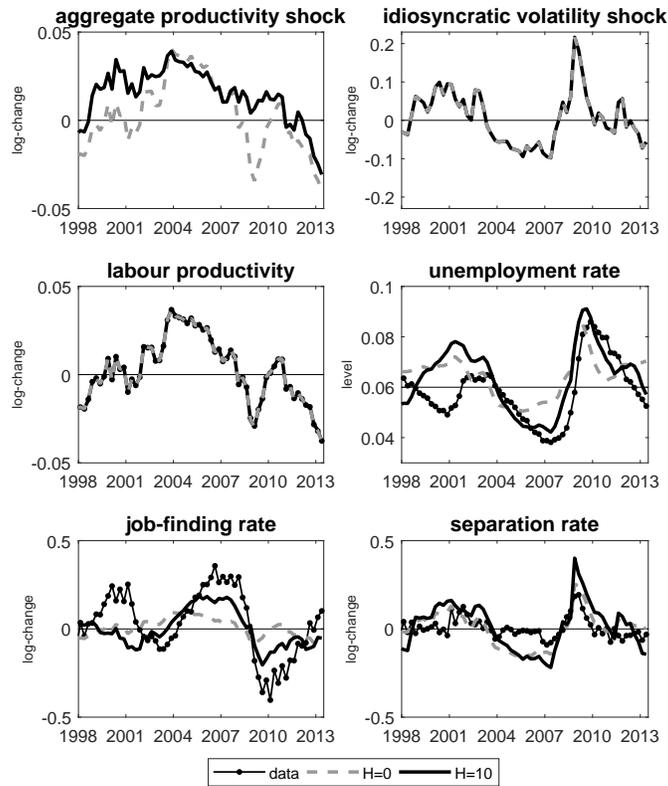
6   Comparison with US time series

In this section, I compare the model with the US time series from 1998 to 2013. Uncertainty shocks seem a good candidate to explain the low unemployment rate in 2007, as well as the decline in unemployment since 2010. In the data, uncertainty was low in the mid 2000s, it peaked in 2008, and it has been falling afterwards. In the model, low uncertainty reduces unemployment.[34]

I use the estimated index of implied volatility, shown in Figure 1 as a proxy for idiosyncratic volatility.[35] The aggregate productivity shock is estimated in each quarter to match labour productivity in the model to the data.

Figure 9 shows the resulting time series for the unemployment rate, the

---

[34]Appendix E shows that productivity shocks alone cannot match the observed unemployment rate after 2004 well.

[35]I use a scaled down version of the log-changes in implied volatility because employment in the model reacts more strongly to the same relative change in idiosyncratic volatility than employment in the data reacts to changes in implied volatility.

Figure 9: Comparison with US time series
Notes. Model with aggregate productivity and uncertainty shocks: Aggregate productivity is estimated in each quarter to match labour productivity in the model to the data. Log-changes in idiosyncratic volatility are 20% of changes of the implied volatility index estimated in Section 2. Linear trends are removed from US data series.

job-finding rate and the separation rate. The model without creation costs cannot reproduce the low unemployment rate in 2007 and fluctuations of the job-finding rate are too small. When creation cost $H = 10$ is used instead, the model closely tracks the unemployment rate in the 2000s. It can explain a steep fall in the job-finding rate after 2007. But the job-finding rate is still a bit less volatile than in the data, whereas the separation rate is too volatile. In terms of the timing, the unemployment rate and the job-finding rate lead their respective empirical counterparts. This indicates that market tightness reacts more sluggishly in the data than in the model. If the creation cost was

37

made dependent on the aggregate number of positions created in each period as in Fujita and Ramey (2007), the response of aggregate creation to aggregate shocks would become smoother. This could improve the fit of the model.

The model has difficulties to explain the low unemployment rate in the run-up of the 2001 recession. The reason is that observed volatility was above trend and labour productivity was close to the trend. Then, the model predicts the unemployment rate to be above trend.

7   Conclusion

In this paper, I have analysed the effects of uncertainty on labour market variables in a model with frictional labour markets and costs of creating positions. I demonstrated that in the absence of creation costs, uncertainty shocks hardly affect the job-finding rate. Then, higher uncertainty lowers employment solely by increasing the separation rate. In the data, however, uncertainty shocks reduce employment also by lowering the job-finding rate. When costs of creating positions are added to the model, the job-finding rate responds negatively to higher uncertainty, which brings the model closer in line with the data. The full model with idiosyncratic volatility shocks and costs of creating positions can match the observed US unemployment rate from 2004-2013 well. In addition, it can also account for a large share of the steep drop in the job-finding rate after 2007.

The model has also shown that similar shocks can have quantitatively very different responses depending on when the shock hits. For instance, an increase in uncertainty leads to a particularly large drop in employment, when the unemployment rate is already high. This suggests that policy makers should react differently to shocks depending on the state of the economy. Future work could explore the effectiveness of government policy in booms compared to recessions. The state-dependency found in this paper suggests that macroeconomic variables will react more strongly in response to policy interventions in times of higher unemployment.

References

Bloom, Nick, Stephen Bond, and John Van Reenen (2007). Uncertainty and Investment Dynamics, *Review of Economic Studies*, 74: 391-415.

Bloom, Nicholas (2009). The Impact of Uncertainty Shocks, *Econometrica*, 77(3): 623-685.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry (2018). Really Uncertain Business Cycles, *Econometrica*, forthcoming.

Coles, Melvyn G., and Ali Moghaddasi Kelishomi (2011). New Business Start-ups and the Business Cycle, *CEPR Discussion Paper*, No. 8588.

Cooper, Russell W., John C. Haltiwanger, and Jonathan Willis (2004). Dynamics of Labor Demand: Evidence from Plant-level Observations and Aggregate Implications, *NBER Working Paper*, No. 10297.

Cooper, Russell, John Haltiwanger, and Jonathan L. Willis (2007). Search frictions: Matching aggregate and establishment observations, *Journal of Monetary Economics*, 54: 56-78.

Costain, James S., and Michael Reiter (2008). Business cycles, unemployment insurance, and the calibration of matching models, *Journal of Economic Dynamics and Control*, 2008(4): 1120-1155.

Davis, Stephen J., R. Jason Faberman, and John Haltiwanger (2012). Labor market flows in the cross section and over time, *Journal of Monetary Economics*, 59: 1:18.

Davis, Stephen J., and John Haltiwanger (1992). Gross Job Creation, Gross Job Destruction, and Employment Reallocation, *Quarterly Journal of Economics*, 107(3): 819-863.

den Haan, Wouter J., Garey Ramey, and Joel Watson (2000), Job Destruction and Propagation of Shocks, *American Economic Review*, 90(3): 482-498.

Elsby, Michael W. L., and Ryan Michaels (2013). Marginal Jobs, Heterogeneous Firms, and Unemployment Flows, *American Economic Journal: Macroeconomics*, 5(1): 1-48.

Fujita, Shigeru, and Makoto Nakajima (2016). Worker Flows and Job Flows: A Quantitative Investigation, *Review of Economic Dynamics*, 22: 1-20.

Fujita, Shigeru, and Garey Ramey (2007). Job Matching and Propagation, *Journal of Economic Dynamics and Control*, 31(11): 3671-3698.

Guglielminetti, Elisa (2016). The Labor Market Channel of Macroeconomic Uncertainty, *Bank of Italy Working Paper*, No. 1068.

Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng (2015). Measuring Uncertainty, *American Economic Review* , 105(3): 1177-1216.

Hall, Robert E. (2005). Employment Fluctuations with Equilibrium Wage Stickiness, *American Economic Review*, 95(1): 50-65.

Kaas, Leo, and Philipp Kircher (2015). Efficient Firm Dynamics in a Frictional Labor Market, *American Economic Review*, 105(10): 3030-3060.

Khan, Aubhik, and Julia K. Thomas (2013). Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity, *Journal of Political Economy*, 121(6): 1055-1107.

Krusell, Per, and Anthony Smith (1998). Income and Wealth Heterogeneity in the Macroeconomy, *Journal of Political Economy*, 106(5): 867-896.

Leahy, John V., and Toni M. Whited (1997). The Effect of Uncertainty on Investment: Some Stylized Facts, *Journal of Money, Credit and Banking*, 28(1), 64-83.

Mecikovsky, Ariel M., and Matthias Meier (2017). Do Plants Freeze upon Uncertainty Shocks?, https://sites.google.com/site/matthias1meier1.

Michaillat, Pascal (2014). A Theory of Countercyclical Government Multiplier, *American Economic Journal: Macroeconomics*, 6(1): 190-217.

Mortensen, Dale T., and Christopher A. Pissarides (1994). Job Creation and Job Destruction in the Theory of Unemployment, *Review of Economic Studies*, 61(3): 397-415.

Petrongolo, Barbara, and Christopher A. Pissarides (2001). Looking into the Black Box: A Survey of the Matching Function, *Journal of Economic Literature*, 39(2): 390-431.

Pissarides, Christopher A. (2009). The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?, *Econometrica*, 77(5): 1339-1369.

Postel-Vinay, Fabien, and Jean-Marc Robin (2002). Equilibrium Wage Dispersion with Worker and Employer Heterogeneity, *Econometrica*, 70(6): 2295-2350.

Schaal, Edouard (2017). Uncertainty and Unemployment, *Econometrica*, 85(6), 1675-1721.

Shimer, Robert (2005). The Cyclical Behavior of Equilibrium Unemployment and Vacancies, *American Economic Review*, 95(1): 25-49.

Shimer, Robert (2012). Reassessing the ins and outs of unemployment, *Review of Economic Dynamics*, 15: 127-148.

Stole, Lars A., and Jeffrey Zwiebel (1996). Intra-firm Bargaining under Non-Binding Contracts, *Review of Economic Studies*, 63(3): 375-410.

van Ours, Jan, and Geert Ridder (1992). Vacancies and the Recruitment of New Employees, *Journal of Labor Economics*, 10(2): 138-155.

Young, Eric (2010). Solving the incomplete markets model with aggregate uncertainty using the Krusell-Smith algorithm and non-stochastic simulations, *Journal of Economic Dynamics and Control*, 34(1): 36-41.

Appendices

A   Empirical part

A.1   Data description

The following time series are used for the estimation of the VAR and when comparing the model to the US data.

Labour productivity. $\frac{Y}{L}$ is obtained from the series of seasonally adjusted nonfarm output (PRS85006043) and employment (PRS85006013) from the U.S. Department of Labor: Bureau of Labor Statistics (BLS).

Index of implied volatilities. Implied volatility data from Option-Metrics' Ivy DB database, available from 1996 until the second quarter of 2013, is used. Call options, whose underlying security is not an index, with a maturity of 30 days are used. Quarterly averages of log implied volatility are taken for each of the underlying securities. This gives between 1,758 and 3,836 observations per quarter. Then coefficients of dummy variables for each quarter are estimated, allowing for fixed effects for the securities. These dummy variables, which capture the average implied volatility, form the index of implied volatilities used in this paper.

Unemployment rate, job-finding rate, and separation rate. They are constructed using data from the Current Population Survey (CPS). The seasonally adjusted unemployment rate is constructed by the BLS (LNS14000000). The job-finding rate and separation rate are calculated using the level of unemployed (LNS13000000) and the number of unemployed for less than 5 weeks (LNS13008396). I account for the redesignment of the unemployment duration question in the CPS, which led to a discontinuous drop of short term unemployment in 1994, as discussed in Shimer (2012). Therefore, the number of newly unemployed is multiplied by the constant 1.1 from 1994. I adjust for time aggregation as emphasized by Shimer (2012) to account for the possibility that workers experience more than one switch between unemployment and

employment per month. The monthly job-finding probability is given by[36]

$$F_t = 1 - \frac{u_{t+1} - u_{t+1}^S}{u_{t+1}}, \tag{A.1}$$

which gives the job-finding rate in a continuous time environment:

$$f_t = -\log\left(1 - F_t\right). \tag{A.2}$$

The separation rate in continous time s_t is then implicitly given by

$$u_{t+1} = \frac{1 - \exp\left(-f_t - s_t\right)s_t}{f_t + s_t}l_t + \exp\left(-f_t - s_t\right)u_t. \tag{A.3}$$

Quarterly averages of the monthly unemployment rate, job-finding rate and separation rate are taken.

A.2   Approximation of changes in unemployment

The law of motion of the unemployment rate is given by

$$u_{t+1} = s_t\left(1 - u_t\right) + \left(1 - f_t\right)u_t. \tag{A.4}$$

If the separation rate s_t and the job-finding rate f_t are constant, the unemployment rate converges to its steady state value

$$\bar{u} = \frac{\bar{s}}{\bar{s} + \bar{f}}. \tag{A.5}$$

The rate of convergence is determined by the sum of $\bar{s}$ and $\bar{f}$:

$$u_{t+1} - \bar{u} = \left[1 - \left(\bar{s} + \bar{f}\right)\right]\left(u_t - \bar{u}\right). \tag{A.6}$$

For example, when the monthly separation rate and job-finding rate sum to 50%, which is roughly the case for the US, any gap between the initial un-

---

[36]In what follows, u_t denotes unemployment, u_t^S denotes unemployment less than 5 weeks, and l_t denotes employment.

employment rate and the steady state value $\bar{u}$ is reduced by 87.5% within a quarter. Therefore, the unemployment rate can be approximated well by only using the current separation rate and job-finding rate, as observed for example by Shimer (2012):

$$u_{t+1} \approx \frac{s_t}{s_t + f_t}. \tag{A.7}$$

Log-linearizing this relationship around $\bar{u}$, $\bar{s}$, and $\bar{f}$, yields

$$
\begin{aligned}
\log \frac{u_{t+1}}{\bar{u}} &\approx \log \frac{s_t}{\bar{s}} - \frac{1}{s_t + f_t} \left(s_t \log \frac{s_t}{\bar{s}} + f_t \log \frac{f_t}{\bar{f}} \right) \\
&= \frac{f_t}{s_t + f_t} \left(\log \frac{s_t}{\bar{s}} - \log \frac{f_t}{\bar{f}} \right) \\
&= (1 - \bar{u}) \left(\log \frac{s_t}{\bar{s}} - \log \frac{f_t}{\bar{f}} \right),
\end{aligned}
\tag{A.8}
$$

where the last equality uses the definition of the steady state unemployment rate.

An analogous relationship can be derived for employment:

$$
\begin{aligned}
n_{t+1} &\approx \frac{f_t}{s_t + f_t}, \tag{A.9} \\
\log \frac{n_{t+1}}{\bar{n}} &\approx (1 - \bar{n}) \left(-\log \frac{s_t}{\bar{s}} + \log \frac{f_t}{\bar{f}} \right) \\
&= \bar{u} \left(-\log \frac{s_t}{\bar{s}} + \log \frac{f_t}{\bar{f}} \right). \tag{A.10}
\end{aligned}
$$

This states that relative changes of employment are proportional to the unemployment rate, when the separation rate and the job-finding rate change by the same relative amount. In particular, same relative changes of market tightness have bigger effects on employment when the unemployment rate is high. This is due to the convexity of the quasi labour supply function given by the matching function, as emphasized by Michaillat (2014).
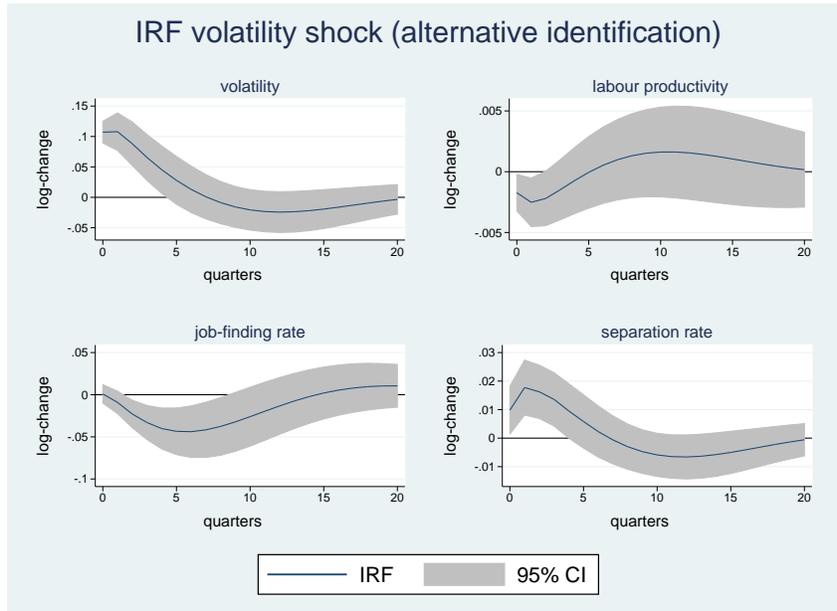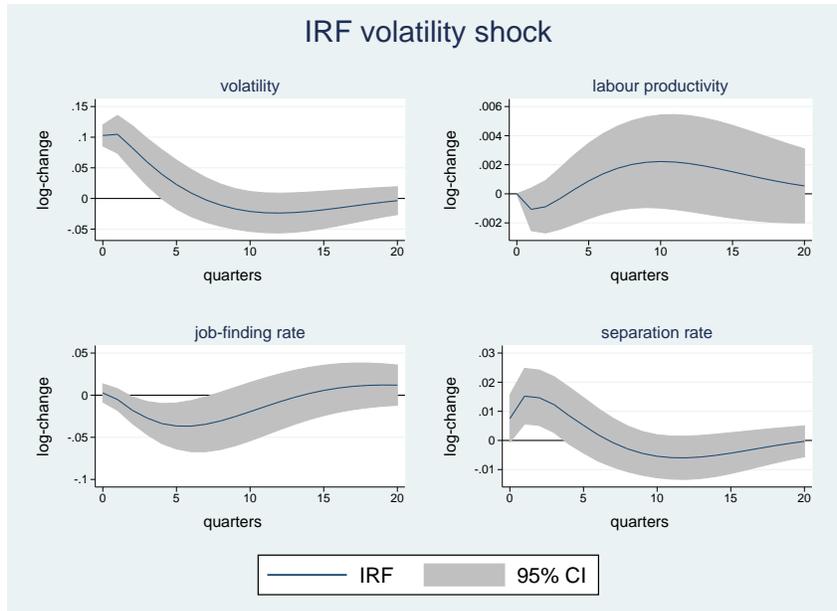
Figure A.1: Impulse response functions after a volatility shock
Notes. An alternative identification, in which volatility is ordered first in the Cholesky decomposition is used.

A.3 Robustness checks

This section shows that the impulse response function in Figure 2 is robust to alternative identification, to an alternative volatility measure, and to alternative detrending of the variables. In Figure A.1, the volatility index is ordered first in the Cholesky decomposition, which implies that labour productivity can react instantaneously to volatility shocks. Figure A.2 uses the volatility index VIX instead of the constructed index based on single stocks. Figures A.3 and A.4 use the Hodrick-Prescott filter to remove trends instead of removing a linear trend in the baseline. Figure A.3 uses the smoothing parameter $\lambda = 100,000$ as preferred by Shimer (2005), whereas Figure A.4 uses the usual value of $\lambda = 1,600$ for quarterly data.
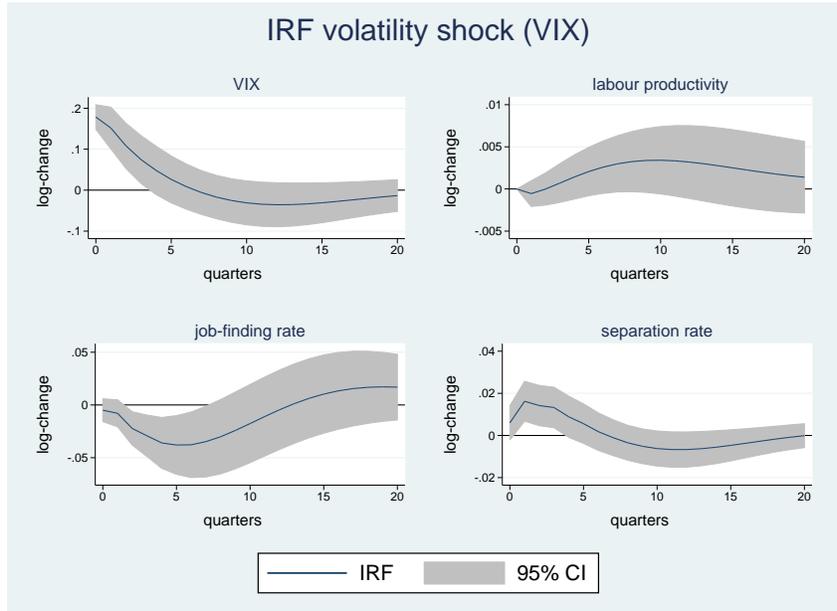
Figure A.2: Impulse response functions after a volatility shock
Notes. Instead of the volatility measure constructed from single stocks, the VIX is used as the volatility measure.

Figure A.3: Impulse response functions after a volatility shock
Notes. Instead of detrending using a linear trend, the series are detrended using the HP-filter with parameter $\lambda = 100,000$.
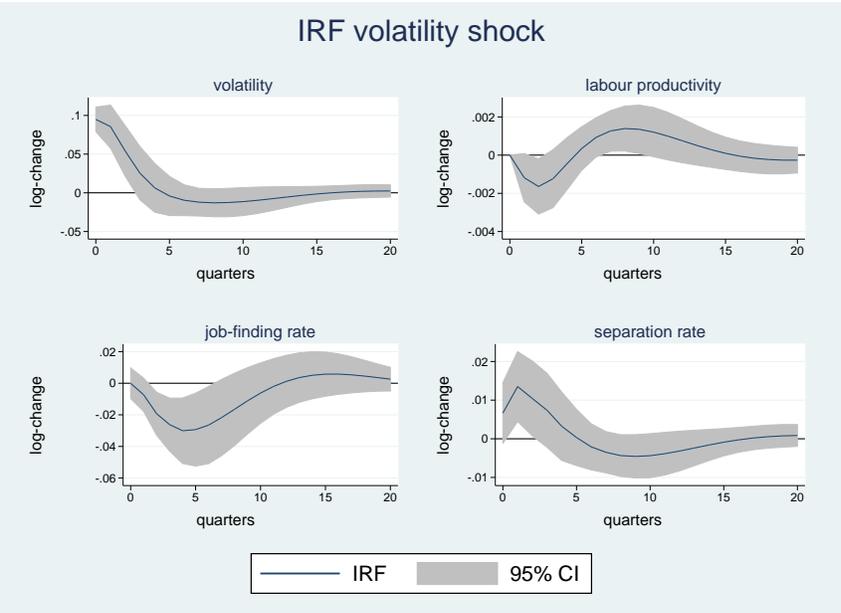
Figure A.4: Impulse response functions after a volatility shock
Notes. Instead of detrending using a linear trend, the series are detrended using the HP-filter
with parameter $\lambda = 1,600$.

B Comparative statics

B.1 Sensitivity analysis

The purpose of this section is to show that the main result of this paper is not sensitive to the choice of parameters. In the absence of creation costs ($H = 0$), the job-finding rate hardly reacts to changes in idiosyncratic volatility, and as the creation cost increases, the falling job-finding rate contributes increasingly to the fall in employment. This is shown in the lower right subplot of the comparative statics graphs in this section.

In the following exercises unless specified otherwise, the three parameters that are calibrated ($\bar{\sigma}_x$, ω_3, and the mass of firms) are recalibrated such that the targets in steady state (unemployment rate, job-finding rate, filling cost relative to average wage) are met exactly.

Destruction cost instead of creation cost. In this exercise, there is no cost of creation, but a cost H has to be paid, whenever a position is shut down voluntarily. The results in Figure B.1 are almost identical to the cost paid upon creation. The reason is that firms are forward looking and take into account that they have to pay the destruction cost when they want to close the position in the future. As a result, the marginal value of a position shifts down and the three relevant values defining the cutoffs in equation (10) become $\left(-H, -H + (1 - \lambda)\frac{c}{q}, (1 - \lambda)\frac{c}{q}\right)$ instead of $\left(0, (1 - \lambda)\frac{c}{q}, H + (1 - \lambda)\frac{c}{q}\right)$.

Same absolute increase of idiosyncratic volatility. In this paper, shocks to uncertainty are interpreted as relative changes of idiosyncratic volatility. The calibration strategy results in steady state volatility $\bar{\sigma}_x$ to be increasing in creation cost, because a higher creation cost makes firms more reluctant to fire. In order to meet the target for the separation rate, volatlity has to be increased. This implies that a 10% increase of volatility increases volatility by a larger absolute amount when H is high, which explains the stronger response of employment. Figure B.2 shows how the results differ when volatility is increased by 0.02, independent of H. Then, the response of employment becomes almost flat, but the contribution of the job-finding rate to its change is hardly affected. Both the separation rate and the job-finding rate react less
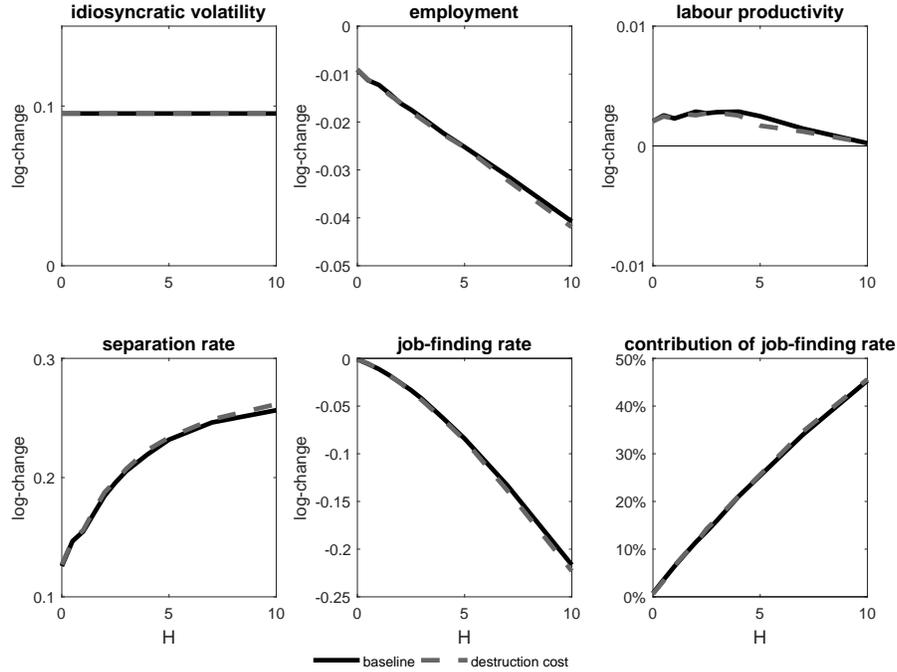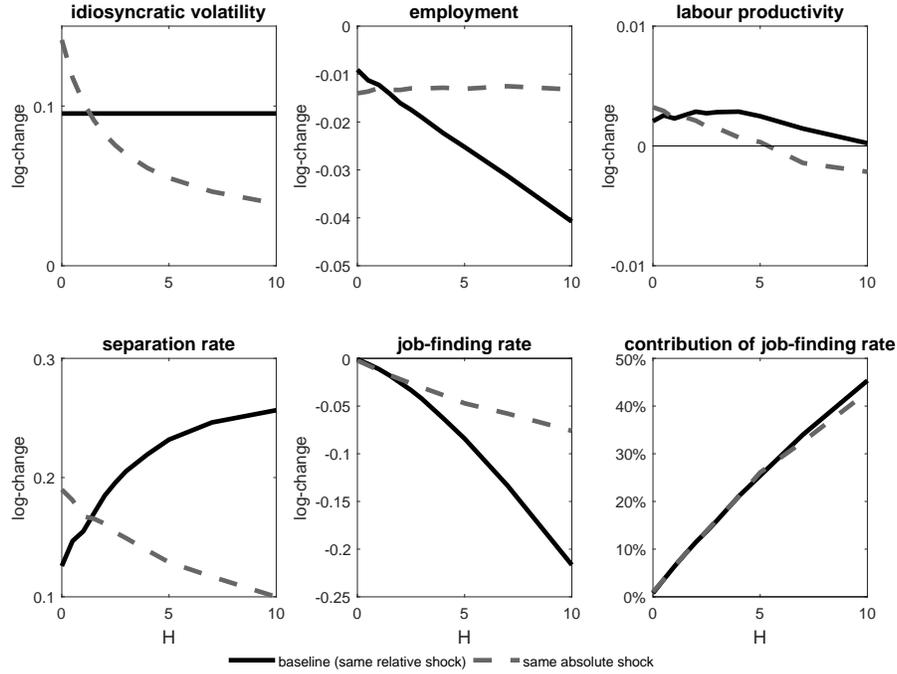
Figure B.1: Sensitivity with respect to destruction cost

Notes. Responses after a permanent 10% increase of idiosyncratic volatility for varying H. In the alternative solution, the cost is not paid upon creation but upon destruction of a position.

for large values of H.

Cyclical creation cost. Pissarides (2009) has shown that fluctuations of unemployment in the standard search and matching model are amplified if the total vacancy posting cost consists of a payment that is independent of market tightness. Without this assumption, higher market tightness in a boom makes it more costly for firms to fill their positions, which dampens the increase of vacancy posting. In my model, the cost of creating a position is assumed to be constant. This is not directly comparable with the assumption in Pissarides (2009) because the cost does not have to be paid when an empty position needs to be filled. The cost of finding workers in my model depends on market tightness as in a standard search and matching model. It is not clear whether the cost of creating (or destroying) positions should be higher when labour market tightness is high. If that were the case, higher costs
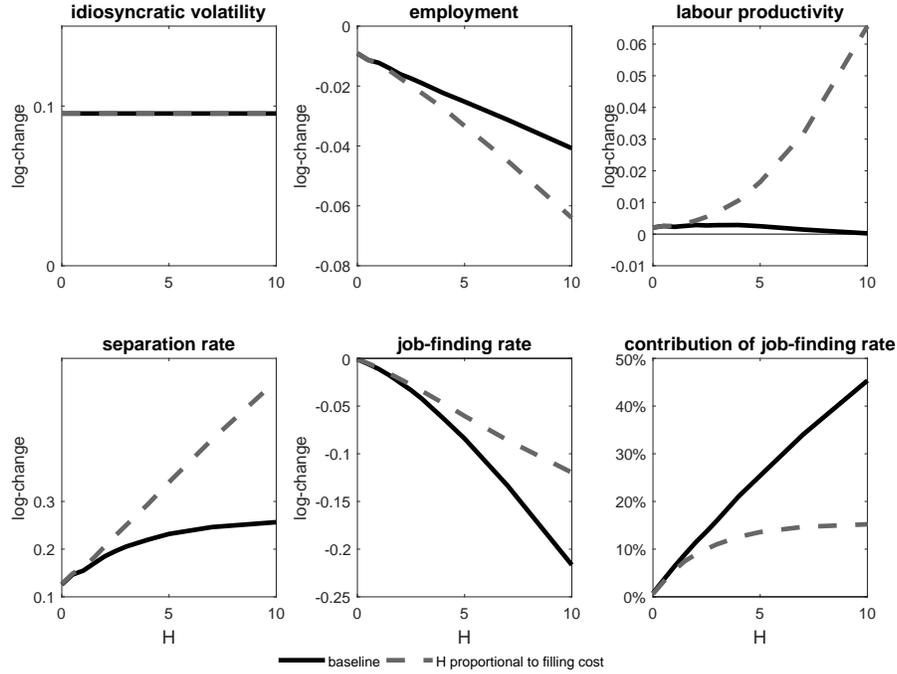
Figure B.2: Sensitivity with respect to size of volatility shock
Notes. Responses after a permanent 10% increase of idiosyncratic volatility for varying H. In the alternative solution, idiosyncratic volatility increases by 0.02 for all H.

of creating positions would dampen booms in addition to the higher filling cost. However, this does not affect the result that the job-finding rate does not move in response to uncertainty shocks when $H = 0$, because if market tightness does not react, it is irrelevant whether the cost of creating positions depends on market tightness. Figure B.3 shows the results when the cost of creating positions is proportional to the filling cost $\frac{c}{q(\theta)}$.[37] The result is that for high values of H, the cost of creation falls relative to the baseline scenario. Consequently, firms become more willing to shut down positions, because it is cheaper to create them in the future and the separation rate increases. For the same reason, the job-finding rate is reduced because firms are less willing to fill their empty positions. As a result, the job-finding rate contributes less to the

---

[37]Remember that in the initial steady state without aggregate uncertainty, $\frac{c}{q(\theta)}$ is normalized to be 1. Now, H changes in proportion to changes in $\frac{c}{q(\theta)}$ in the new steady state.
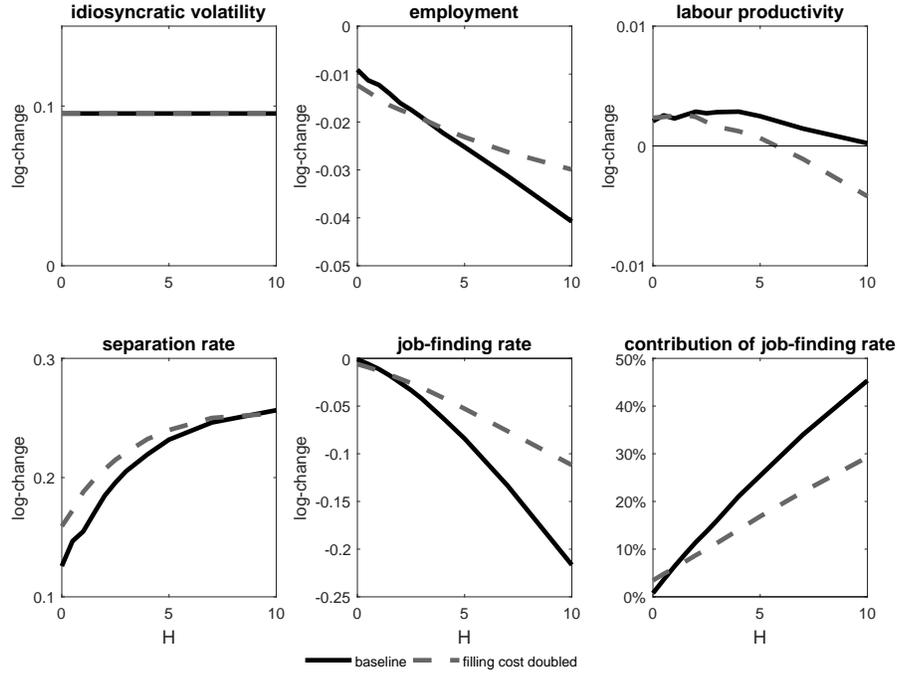
Figure B.3: Sensitivity with respect to cyclicality of H
Notes. Responses after a permanent 10% increase of idiosyncratic volatility for varying H.
In the alternative solution, the creation cost H is proportional to the filling cost $\frac{c}{q}$.

fall in employment, which becomes larger. Both lower employment and the reduced importance of the wait-and-see effect lead to an increase in measured labour productivity.

Filling cost. If it is more costly to fill positions, the wait-and-see effect becomes stronger even in the absence of a creation cost. Figure B.4 shows that even doubling the target of the filling cost relative to average wages increases the contribution of the job-finding rate to less than 5% when $H = 0$. For larger values of H, the job-finding rate actually falls less than with the lower filling cost. The reason it that firms become more sensitive to changes in market tightness, because the filling costs are more important for them. This increases labour demand more in response to falling market tightness such that in general equilibrium the job-finding rate falls less.

Exogenous turnover. Figure B.5 shows the solution for alternative val-
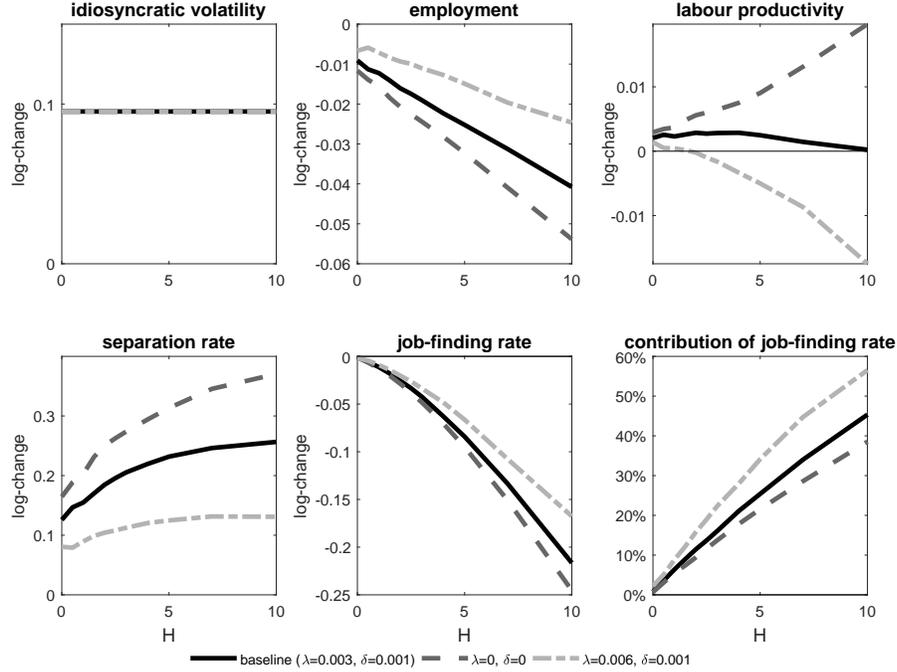
51

Figure B.4: Sensitivity with respect to filling cost

Notes. Responses after a permanent 10% increase of idiosyncratic volatility for varying H. In the alternative solution, the target for the filling cost is doubled relative to the baseline scenario. The creation cost remains unchanged, such that $H = 10$ now corresponds to the creation cost being 5 times the filling cost in steady state.

ues of the rate of obsolescence δ and quits λ. If both parameters are set to 0 as in Elsby and Michaels (2013), both the job-finding rate and the separation rate react more to higher uncertainty, and the job-finding rate contributes slightly less to changes in employment. In constrast, if λ was doubled, both labour market transition rates react less and the job-finding rate contributes more. Note that the calibration strategy ensures, that the separation rate is the same in steady state. This means that there are more *endogenous* separations when δ and λ are small. Then, a shock that increases endogenous separations by the same relative amount, leads to a bigger increase of total separations.

Alternative wage equation. Figure B.6 shows the results for alternative parameters in the wage equation (15). The dashed solution halves the parameter ω_2. This makes the wage less responsive to changes in market tightness.
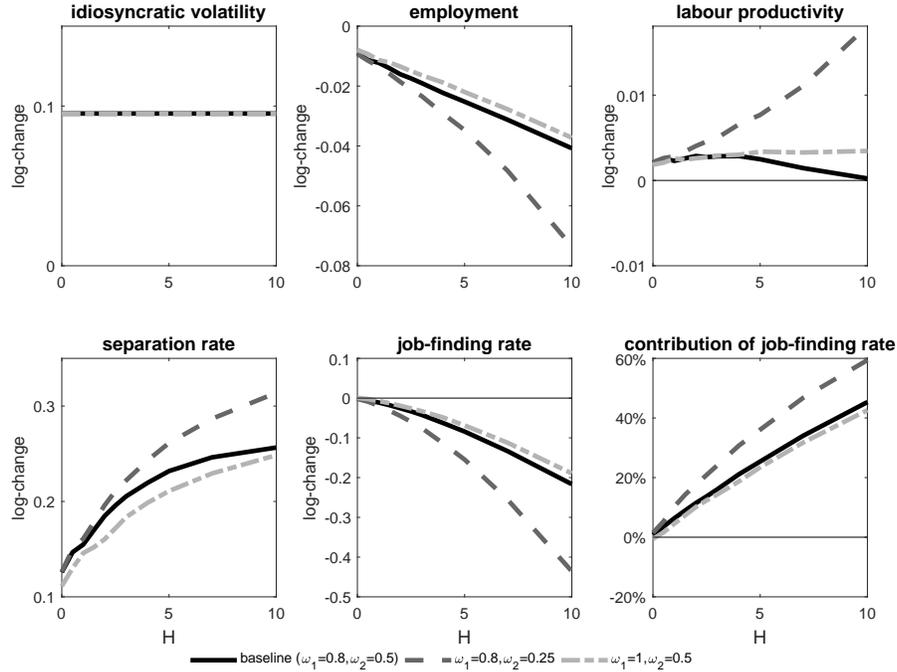
Figure B.5: Sensitivity with respect to exogenous turnover

Notes. Responses after a permanent 10% increase of idiosyncratic volatility for varying H. The dashed solution solves the model without depreciation of positions and without quits. The dash-dotted solution solves the model for a doubled rate of quits.

Therefore, market tightness and the job-finding rate need to fall more for high values of H to ensure market clearing. In the dash-dotted solution, the parameter ω_1 is set to 1, which means that the wage does not depend on firm size. Increasing only ω_1 means that firms become more profitable as if they had a permanently higher productivity. This scales up their demand for labour and the average firm size. A corresponding decrease of the mass of firms could perfectly offset this, and the results would be identical to the baseline scenario. The only reason that they differ is that the wage coefficient ω_3 is recalibrated to meet the target for the filling cost relative to the wage. This has only minor consequences for the contribution of the job-finding rate to changes in employment.

Curvature of production function. In the literature of multi-worker firm models with decreasing returns to scale, a wide range of parameter values
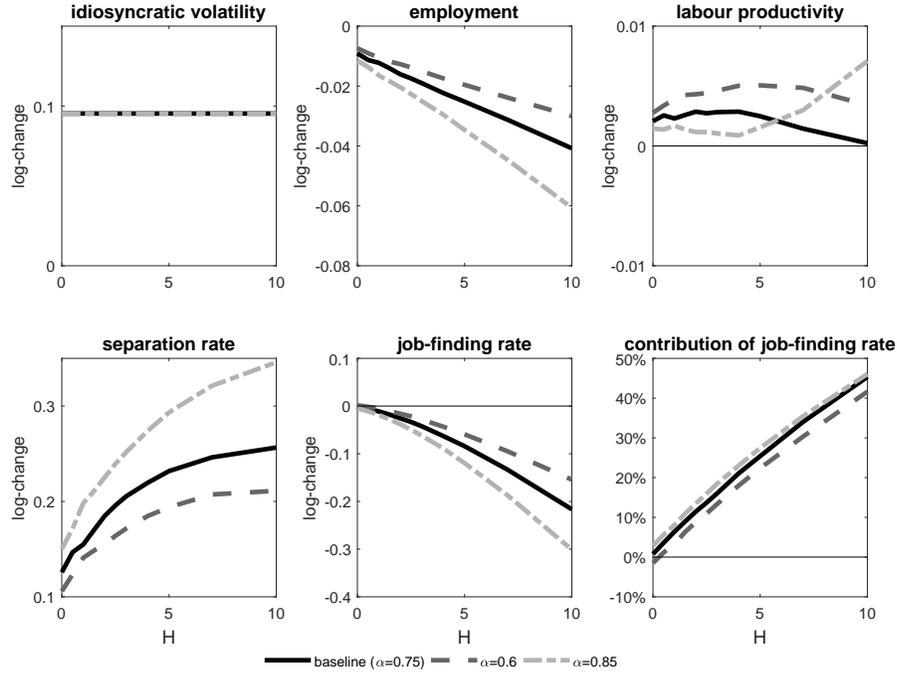
Figure B.6: Sensitivity with respect to wage equation

Notes. Responses after a permanent 10% increase of idiosyncratic volatility for varying H. The dashed solution uses a wage equation that is less sensitive to changes in market tightness, and the dash-dotted solution makes the wage independent of firm size.

for the elasticity of the production or revenue function has been used. In the baseline scenario, I use an intermediate value of $\alpha = 0.75$. Figure B.7 confirms that lowering it to 0.6 as in Elsby and Michaels (2013) or increasing it to 0.85 as in Schaal (2017) hardly affects the contribution of the job-finding rate to changes in employment. What α does affect is the response of employment, because it determines how much the marginal profit of the firm increases with fewer workers. When α is low, firms adjust their desired number of workers by less in response to shocks than for higher values of α.

Arrival rate of idiosyncratic shock. Figure B.8 shows the results for different values of the arrival rate of the idiosyncratic shock. Because the idiosyncratic productivity draws are uncorrelated in the baseline scenario, a higher arrival rate makes firms less responsive to their current level of productivity. Then, they also react less to changes in idiosyncratic volatility. This
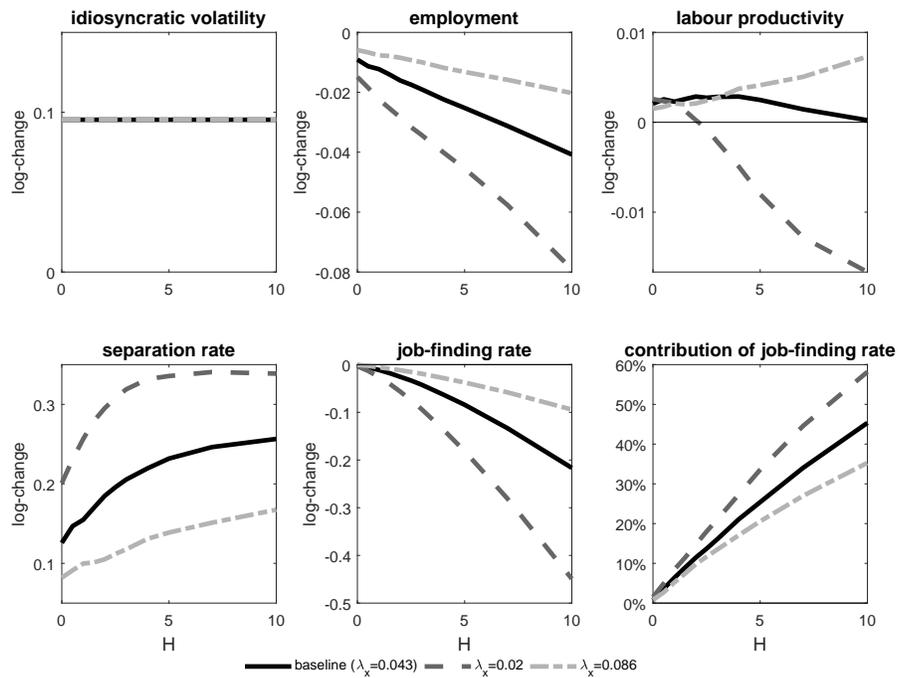
Figure B.7: Sensitivity with respect to α

Notes. Responses after a permanent 10% increase of idiosyncratic volatility for varying H. The alternative solutions use different values of curvature in the production function.

primarily affects changes in employment, but it also reduces the contribution of the job-finding rate.

Distribution of idiosyncratic shock. Figure B.9 shows that the contribution of the job-finding rate to changes in employment is not sensitive to the assumed distribution of idiosyncratic productivity. The dashed curves show the result when a Pareto distribution instead of the log-normal distribution is used as in Elsby and Michaels (2013). The dash-dotted curves assume that idiosyncratic productivity changes in each period, but it is autocorrelated as in Schaal (2017). In both cases, employment reacts less to uncertainty shocks, but the contribution of the job-finding rate increases slightly.

Figure B.8: Sensitivity with respect to arrival rate of shock
Notes. Responses after a permanent 10% increase of idiosyncratic volatility for varying H. In the alternative solutions, the arrival rate of the idiosyncratic shock is varied.
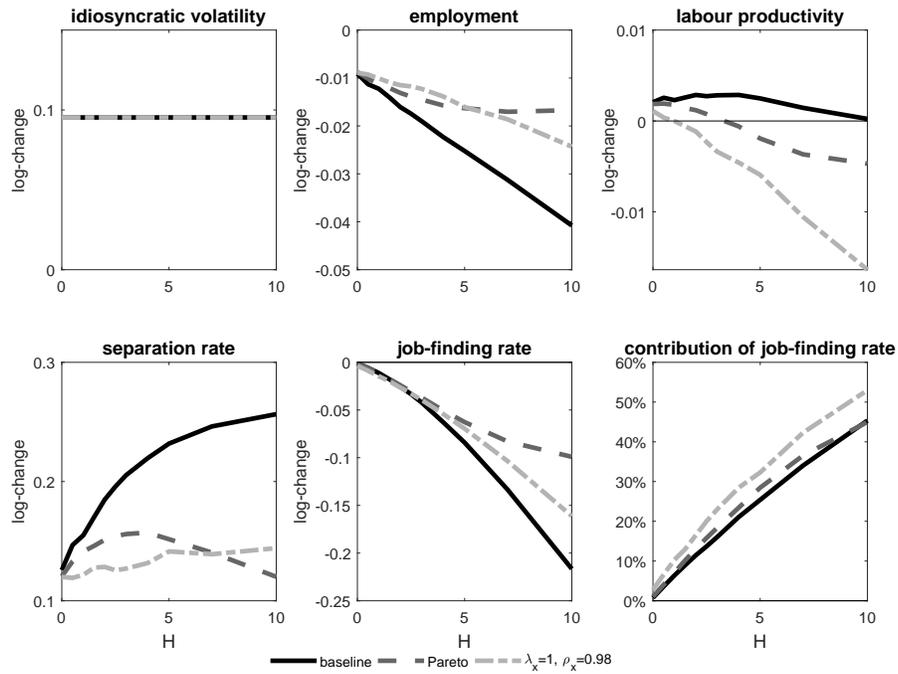
Figure B.9: Sensitivity with respect to idiosyncratic shock

Notes. Responses after a permanent 10% increase of idiosyncratic volatility for varying H. The dashed solution uses a Pareto distribution instead of the log-normal distribution for the idiosyncratic productivity draw. The dash-dotted solution assumes that shocks arrive in each period but that productivity is autocorrelated.

C   Decomposition exercise

C.1   Derivation of change in cutoffs for the wait-and-see effect

The wait-and-see effect measures the change in aggregate hiring and firing that is due to firms adjusting their cutoffs in response to increased uncertainty. It is not useful, however, to just look at the change of cutoffs conditional on idiosyncratic productivity x, like the amount that $n_{LB}(x)$ shifts. This is because firms with the same x are not directly comparable as they are at different percentiles of the cross-sectional distribution when volatility changes. For example, take a firm, whose idiosyncratic productivity x is one standard deviation below the mean. If the standard deviation is doubled, a firm with this same x is only half of a standard deviation below the mean. Consequently, it is more likely that this particular firm gets a new productivity draw that lies below x in the presence of higher uncertainty. Therefore, I will compare how the cutoffs of firms at the same percentile of the cross-sectional distribution change.

It is not possible to just compare the change in cutoffs, because firms at the same percentile have a different value of x depending on the standard deviation of the cross-sectional distribution. In order to account for this mean effect, I will consider how the firm's cutoff changes relative to the "perfect insurance" solution. In this context, "perfect insurance" means the assumption that a firm can always sell an existing position or existing workers to other firms and thereby transfer the position or worker to this other firm. In equilibrium, other firms would be willing to pay exactly the cost that they would otherwise incur to create a position or to hire a worker, respectively. Hence, in terms of the marginal value of a position, firms are perfectly insured against idiosyncratic shocks. If they become less productive, they can sell some of their existing positions and workers to other firms until the marginal value of a position equals the price at which it could be sold. Then equation (8) simplifies, and

the number of positions is determined by

$$\omega_1 \alpha z x n^{\alpha-1} - (\omega_2 \theta + \omega_3) + \beta (1 - \delta) H + \beta (1 - \delta) (1 - \lambda) \mathbb{E} \left[\frac{c}{q (\theta (s'))} | s \right]$$
$$= H + \frac{c}{q (\theta (s))}. \quad \text{(C.1)}$$

This states that the marginal product minus the marginal wage plus the discounted expected marginal value must be equal to the cost of creating and filling a position. The marginal value of a position created and filled today in the next period consists of the value of the position, $(1 - \delta) H$, taking into account the exogenous destruction of a share δ of positions, and the value of the worker filling it, $(1 - \delta) (1 - \lambda) \mathbb{E} \left[\frac{c}{q(\theta(s'))} | s \right]$, taking into account that a share λ of workers in not exogenously destroyed positions will leave the firm. Note that the only risk remaining for the firm is aggregate risk, which determines market tightness and hence the filling rate q'. For the remainder of this section, I will abstract from the effects expected changes in market tightness have on firms, and assume that $q' = q$. Then one can solve for the number of positions as a function of idiosyncratic productivity:

$$\log \tilde{n}_{LB} (x) = \frac{\log x + \log (\omega_1 \alpha z)}{1 - \alpha}$$
$$- \frac{\log \left\{ \omega_2 \theta + \omega_3 + [1 - \beta (1 - \delta)] H + [1 - \beta (1 - \delta) (1 - \lambda)] \frac{c}{q} \right\}}{1 - \alpha}. \quad \text{(C.2)}$$

This shows that the number of positions $\tilde{n}_{LB} (x)$ is log-linear in idiosyncratic productivity in the "perfect insurance" case. We can compare this solution to the creation cutoff $n_{LB} (x)$ in the model. In general, $n_{LB} (x) \leq \tilde{n}_{LB} (x)$ holds, because the continuation value in the model with idiosyncratic risk never exceeds the continuation value in the perfect insurance model. The difference between the creation cutoff and the perfect insurance number of positions can be interpreted as the amount of positions that a firm creates less in the presence of idiosyncratic risk, because it fears that the marginal value of the newly created positions falls in the future. One can show that $\lim_{x \to 0} (n_{LB} (x) - \tilde{n}_{LB} (x)) = 0$

holds, because for low values of idiosyncratic productivity it becomes increasingly unlikely that an even lower value of productivity is drawn in the next period. In particular, if there is a minimum value in the support of the distribution of idiosyncratic productivity, $n_{LB}(x_{\min}) = \tilde{n}_{LB}(x_{\min})$. Then, a firm knows for sure that productivity can't fall in the next period, and that the continuation value is equal to the creation and filling cost as in the perfect insurance economy. Figure C.1 draws the perfect insurance cutoff in addition to the creation cutoff for $H = 5$ and $H = 0$. Note that the difference between the two is larger for $H = 5$, because the continuation value could potentially fall by more. In the limiting case of no creation cost and no search frictions ($c = 0$), the two cutoffs would coincide as well.
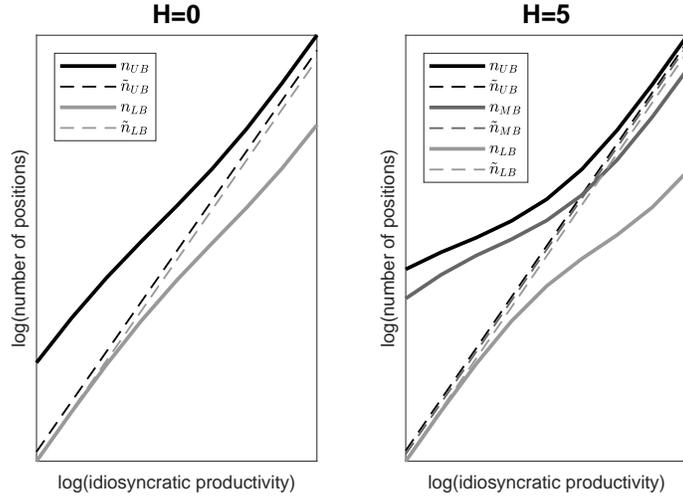
Figure C.1: Optimal cutoffs (solid) and "perfect-insurance" cutoffs (dashed)

Now, I can define analogous measures for the middle cutoff $n_{MB}(x)$, which determines whether a firm wants to fill empty positions, and for the upper cutoff $n_{UB}(x)$, which determines how many workers a firm wants to fire. Ac-

cording to equation (7), the middle cutoff is given by the condition

$$
\begin{aligned}
J\left(n_{MB}\left(x\right),x,s\right) &= \omega_1\alpha zxn_{MB}\left(x\right)^{\alpha-1}-\left(\omega_2\theta+\omega_3\right)-\lambda\frac{c}{q\left(\theta\left(s\right)\right)}\\
&\quad+\beta\left(1-\delta\right)\mathbb{E}\left[J\left(\left(1-\delta\right)n_{MB}\left(x,s\right),x',s'\right)|x,s\right]\\
&= \left(1-\lambda\right)\frac{c}{q\left(\theta\left(s\right)\right)}. \quad\quad\quad\quad\quad\text{(C.3)}
\end{aligned}
$$

I want to compare it to a hypothetical cutoff, $\tilde{n}_{MB}\left(x\right)$, that assumes that the expected marginal continuation value is equal to the current period's marginal value:[38]

$$
\mathbb{E}\left[J\left(\left(1-\delta\right)\tilde{n}_{MB}\left(x\right),x',s'\right)|x,s\right] = J\left(\tilde{n}_{MB}\left(x\right),x,s\right). \quad\quad\text{(C.4)}
$$

Then the cutoff $\tilde{n}_{MB}\left(x\right)$ can be calculated:

$$
\begin{aligned}
\log\tilde{n}_{MB}\left(x\right) = & \frac{\log x+\log\left(\omega_1\alpha z\right)}{1-\alpha}\\
& -\frac{\log\left\{\omega_2\theta+\omega_3+\left[1-\beta\left(1-\delta\right)\left(1-\lambda\right)\right]\frac{c}{q\left(\theta\left(s\right)\right)}\right\}}{1-\alpha}. \quad\text{(C.5)}
\end{aligned}
$$

A comparison of this cutoff with the perfect insurance cutoff in equation (C.2) shows that it is obtained by a parallel upward shift in the $(\log x,\log n)$-space. Note that the cutoff $n_{MB}\left(x\right)$ is above $\tilde{n}_{MB}\left(x\right)$ for low values of productivity and below it for high values of x. The reason is that low-productivity firms are more likely to get a positive productivity shock, which leads to a contin- uation value above $\left(1-\lambda\right)\frac{c}{q}$. Therefore, the expected continuation value for firms subject to idiosyncratic risk is higher than $\left(1-\lambda\right)\frac{c}{q}$ in the definition of $\tilde{n}_{MB}\left(x\right)$. This implies that $n_{MB}\left(x\right)>\tilde{n}_{MB}\left(x\right)$ for low x. In contrast, high productivity firms are more likely to experience a negative productivity shock, that results in a continuation value below $\left(1-\lambda\right)\frac{c}{q}$. Hence, the expected con- tinuation value is below $\left(1-\lambda\right)\frac{c}{q}$ for those firms and $n_{MB}\left(x\right)<\tilde{n}_{MB}\left(x\right)$. Finally, note that in the absence of creation costs the middle cutoffs coincide

---

[38]Note that if both the aggregate and idiosyncratic state remain unchanged, $J\left(\left(1-\delta\right)n_{MB}\left(x\right),x',s'\right)>J\left(n_{MB}\left(x\right),x,s\right)$ holds because of depreciation ($\delta>0$). The underlying assumption for the decomposition exercise is that this bias does not depend much on the size of idiosyncratic volatility.

with the low cutoffs as shown in the right panel of Figure C.1.

The upper cutoff is implicitly defined by equation (4)

$$0 = J\left(n_{UB}\left(x,s\right),x,s\right)$$
$$= \left(1-\lambda\right) \left\{ \begin{array}{c} \omega_1 \alpha z x \left[\left(1-\lambda\right) n_{UB}\left(x,s\right)\right]^{\alpha-1} - \left(\omega_2\theta + \omega_3\right) \\ +\beta\left(1-\delta\right)\mathbb{E}\left[J\left(\left(1-\delta\right)\left(1-\lambda\right) n_{UB}\left(x,s\right),x',s'\right)|x,s\right] \end{array} \right\}. \quad (C.6)$$

Analogous to above, I define $\tilde{n}_{UB}\left(x\right)$ under the assumption that

$$\mathbb{E}\left[J\left(\left(1-\delta\right)\left(1-\lambda\right)\tilde{n}_{UB}\left(x\right),x',s'\right)|x,s\right] = 0. \quad (C.7)$$

Then, the cutoff is given by

$$\log \tilde{n}_{UB}\left(x\right) = \frac{\log x + \log\left(\omega_1\alpha z\right) - \log\left(\omega_2\theta + \omega_3\right)}{1-\alpha} - \log\left(1-\lambda\right). \quad (C.8)$$

This cutoff is parallel to $\tilde{n}_{MB}\left(x\right)$ and $\tilde{n}_{LB}\left(x\right)$ in the $\left(\log x, \log n\right)$-space[39]. In general, $n_{UB}\left(x\right) \geq \tilde{n}_{UB}\left(x\right)$, because the marginal continuation value that is relevant for $n_{UB}\left(x\right)$ is nonnegative, whereas it is 0 by assumption for $\tilde{n}_{UB}\left(x\right)$. The difference between the two cutoffs can be interpreted as the amount of workers that are not fired, and whose positions are not shut down, because the firm has a chance to get a positive productivity shock. The difference is getting smaller for high levels of productivity because it becomes less likely for the firm to become even more productive in the future[40].

Let $d_j^\sigma\left(p\right)$ be the log-difference between cutoffs n_j and $\tilde{n}_j$, evaluated at the $100p^{th}$-percentile, when the standard deviation of idiosyncratic productivity is σ.[41] Denoting the inverse cumulative distribution function of the idiosyncratic

---

[39]In theory, $\tilde{n}_{UB}\left(x\right)$ could be smaller than $\tilde{n}_{MB}\left(x\right)$. In my calibration, λ is small enough relative $\frac{c}{q}$ such that $\tilde{n}_{UB}\left(x\right) > \tilde{n}_{MB}\left(x\right)$ holds.

[40]In general it is not true that $\lim_{x\to\infty}\left(n_{UB}\left(x\right) - \tilde{n}_{UB}\left(x\right)\right) = 0$ holds. This is because in the presence of depreciation and attrition of workers, a firm's size shrinks over time. This increases marginal profits leading to a positive expected continuation value even if $J\left(n_{UB}\left(x_{\max}\right),x,s\right) = 0$ and $x' \leq x_{\max}$.

[41]To keep notation simpler, I use subscript j, whenever an expression applies to all three cutoffs LB, MB, and UB.

shock by $F_{\sigma_x}^{-1}(.)$, the three differences can be written as

$$d_{LB}^\sigma(p) = \log n_{LB}\left(F_{\sigma_x}^{-1}(p)\right) - \log \tilde{n}_{LB}\left(F_{\sigma_x}^{-1}(p)\right) \leq 0, \qquad \text{(C.9)}$$

$$d_{MB}^\sigma(p) = \log n_{MB}\left(F_{\sigma_x}^{-1}(p)\right) - \log \tilde{n}_{MB}\left(F_{\sigma_x}^{-1}(p)\right), \qquad \text{(C.10)}$$

$$d_{UB}^\sigma(p) = \log n_{UB}\left(F_{\sigma_x}^{-1}(p)\right) - \log \tilde{n}_{UB}\left(F_{\sigma_x}^{-1}(p)\right) \geq 0. \qquad \text{(C.11)}$$

C.2   Asymmetry of wait-and-see effect

The shifts of the cutoffs for the wait-and-see effect in Figure 5 are not parallel. First, the effects are hump-shaped. They are smaller for p close to 0 or p close to 1, because then most new productivity shocks lead to the firm being at the lower or upper cutoff, respectively. For example, take a firm in a low percentile of current productivity. When a new shock arrives, it is likely to become significantly more productive. In that case, the firm is going to create positions after the good productivity draw even if it was currently (close to) firing workers. Thus the marginal continuation value of a position is likely to be at its upper bound after a new productivity draw. Increasing uncertainty still increases the expected continuation value by a bit, which leads to an increase of $d_{UB}^\sigma(p)$, but the effect is quantitatively small. Second, the change in cutoffs peaks at lower productivity values for the upper bound and at higher productivity values for the lower bound. The reason for this asymmetry is that $|d_{LB}^\sigma(p)|$ is larger for more productive firms as could be seen in Figure C.1. For those firms a future fall in productivity is more likely, which makes them more cautious in creating positions. Higher uncertainty exacerbates the potential losses and increases $|d_{LB}^\sigma(p)|$. In contrast, low-productivity firms are unlikely to get even less productive in the future. Therefore, higher uncertainty increases $|d_{LB}^\sigma(p)|$ by less. A similar argument can be applied for the upper cutoff: $d_{UB}^\sigma(p)$ is small for very productive firms, because they are unlikely to become even more productive. Therefore, changes in uncertainty affect $d_{UB}^\sigma(p)$ less. Note that this asymmetry increases the effects relative to a symmetric shift of the cutoffs, because there a more high-productivity firms close to the lower cutoff and more low-productivity firms close to the upper cutoff. Third,

the effects are quantitatively stronger in the presence of creation cost. The larger the creation cost, the bigger becomes the difference between the upper and the lower bound of the continuation value, and the more likely it is for a firm to be in the interior region after a shock. This means that potential gains or losses from an increase in uncertainty become larger and firms adjust their cutoffs by more. As a result, the wait-and-see effect becomes more relevant when H is large.

C.3 General equilibrium effect
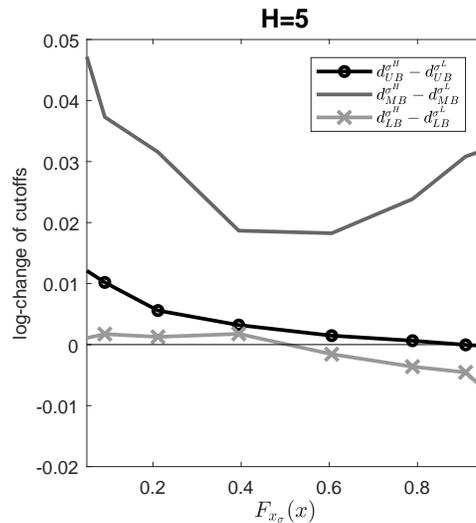
Figure C.2: General equilibrium effect

Notes. The figure draws the change in cutoffs due to a change in market tightness that is necessary to clear markets after a 10% increase in uncertainty. The changes in cutoffs are relative to their "perfect-insurance" counterparts. They are drawn as a function of the percentile in the idiosyncratic productivity distribution ($F_{x_\sigma}(x)$).

In the presence of creation cost, market tightness falls in response to an increase in uncertainty. Figure C.2 shows the change in cutoffs that is due to this change in market tightness after a 10% increase in uncertainty in the comparative statics exercise. It turns out that the biggest change in cutoffs occurs at the middle cutoff n_{MB}. This cutoff is determined by the condition

that a firm is indifferent between filling their empty positions and shutting them down.

Market tightness determines the cost of filling the position, which makes it an important determinant for the cutoff. The creation cutoff n_{LB} instead is determined by the condition that a firm is indifferent between creating and filling a new position or not. When the creation cost is large relative to the filling cost, a lower filling cost makes it more profitable for the firm to create and fill new positions, but the total cost is not affected as much in relative terms as it is for the middle cutoff. Note that the quantitative shift of cutoffs depends on the change in market tightness. For low values of H, the market clearing condition is almost satisfied in partial equilibrium. Then, market tightness does not adjust much, which in turn leads to a small shift of cutoffs.

D   State dependence of IRFs

Figure 8 has shown that the impulse response functions can differ significantly depending on when an uncertainty shock hits the economy. Table D.1 reports the peak change in employment, when the sample of impulse responses is divided into three equally sized parts according to initial unemployment. It shows that uncertainty shocks have a bigger negative impact on employment when employment is already low.

The initial unemployment rate is more important for high values of H. For example when $H = 10$, employment falls by 1.5% when the unemployment rate was in the highest third, and by only 1.1% when the unemployment rate was in the lowest third. One reason, why shocks affect employment more when the unemployment rate is high, is implied by the matching function, as discussed in Michaillat (2014). When the unemployment rate is low, the same relative change of market tightness affects the unemployment rate less. Equation (A.10) in Appendix A.2 shows that changes in log-employment are approximately proportional to the unemployment rate, when the job-finding rate or the separation rate are changed by the same relative amount.

But even if the initial unemployment rate is taken into account, some state

Table D.1: State dependence of impulse response functions

		$d \log n$	
H	n^L	n^M	n^H
0	-0.005	-0.004	-0.004
5	-0.011	-0.009	-0.008
10	-0.015	-0.012	-0.011

Notes. The sample of IRFs is divided into three equal parts according to employment when the volatility shock hits. (n^L, n^M, n^H correspond to the lowest, middle, and highest third, respectively.) Then the average peak changes of log-employment in the subsamples are calculated and reported depending on creation cost H.
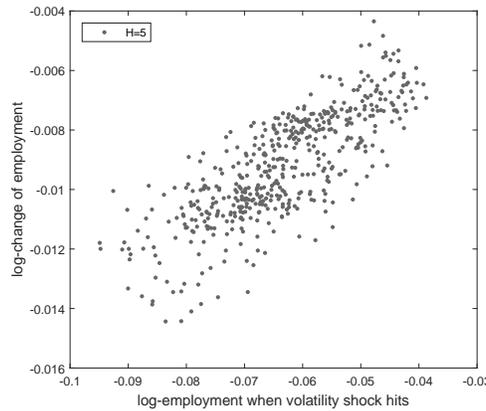
Figure D.1: State dependence of IRFs
Notes. Each dot corresponds to the peak response of log-employment in one impulse response function after a positive uncertainty shock for $H = 5$.

dependency remains. Figure D.1 plots the observed peak responses of employment as a function of initial employment for $H = 5$. It can be seen that the range of possible values conditional on initial employment is still large. These differences are due to characteristics of the cross-sectional distribution that are not captured well by aggregate employment. For example, it could be that in some periods there are more firms close to (one of) the cutoffs, such that they are more responsive to additional shocks. Figure D.1 also gives an indication that the first moment of the distribution is not enough to forecast aggregate outcomes. Appendix F explains the algorithm that I develop to get more ac-

curate forecasts of market tightness without adding many characteristics of the cross-sectional distribution as additional state variables.
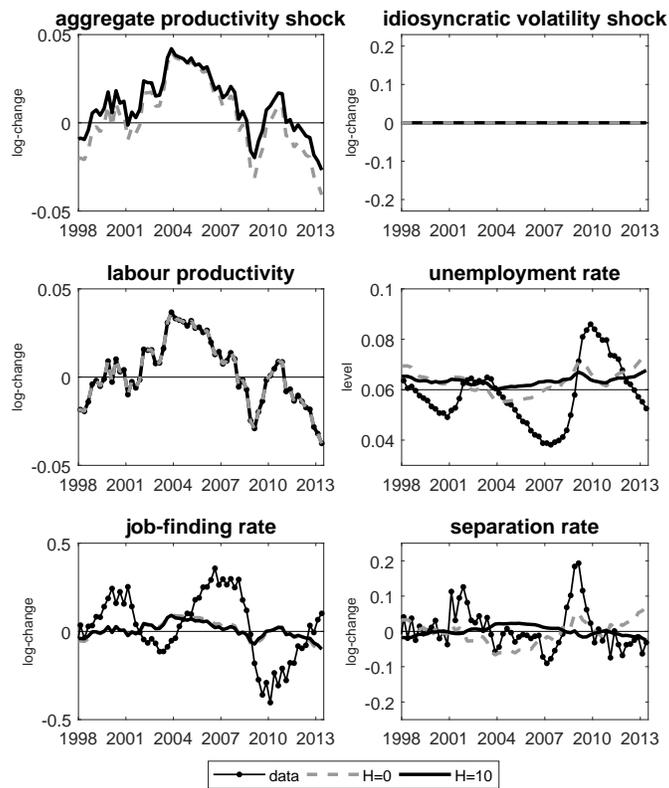
E   Model without uncertainty shocks

In this section, I compare the results of the model with the US time series, when idiosyncratic volatility is constant. I show that productivity shocks alone don't achieve a good fit for the unemployment rate in the great recession. The resulting series of the unemployment rate, the job-finding rate, and the separation rate are drawn in Figure E.1. At first glance, the unemployment rate is not as volatile in the model as in the data.[42]

This shortcoming could be alleviated if the model was recalibrated to produce a larger volatility of employment. In particular, a lower value of ω_2 in the wage equation could be used. Then, the wage reacts less to changes in market tightness. Hence, the marginal profit of positions become more volatile and firms will react more strongly in response. Similarly, Hall (2005) has shown that wage stickiness can help to overcome too little employment volatility in matching models. The results are shown in Figure E.2. Note, however, that even if the unemployment rate reacted more to productivity shocks, it would be difficult to match the observed unemployment rate after 2004. Whereas it was low in 2007, the model would predict the lowest value for unemployment already in 2004. In addition, unemployment increases in the model from 2011 onward, while a steady decline was observed in the data.

---

[42]Too little employment volatility has been documented in standard search and matching models, for example, in Costain and Reiter (2008) and Shimer (2005).

Figure E.1: Model without uncertainty shocks

Notes. Aggregate productivity is estimated in each quarter to match labour productivity in the model to the data. Linear trends are removed from US data series.
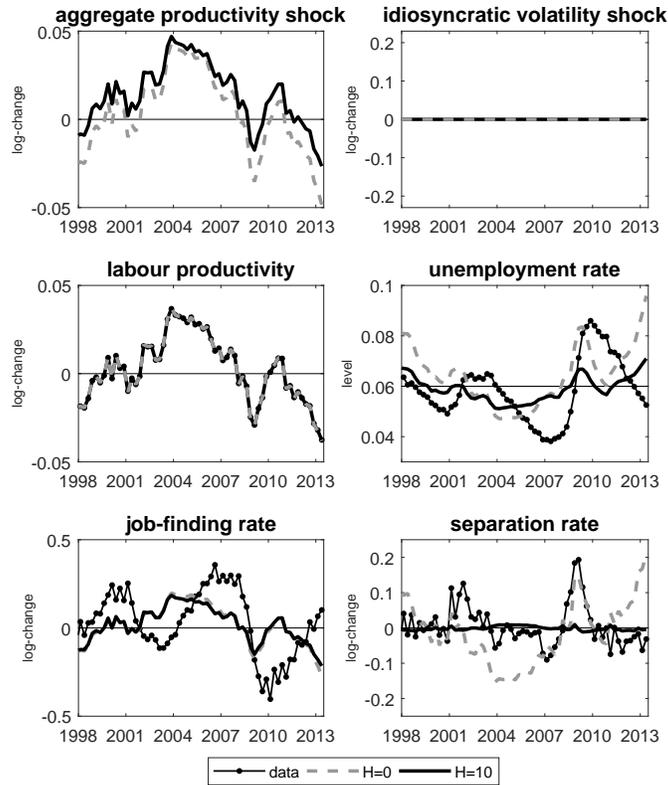
Figure E.2: Model without uncertainty shocks and more wage rigidity

Notes. Aggregate productivity is estimated in each quarter to match labour productivity in the model to the data. Linear trends are removed from US data series. Relative to the baseline calibration, the lower value $\omega_2 = 0.2$ instead of $\omega_2 = 0.5$ is used, which makes the wage less responsive to changes in market tightness and makes the unemployment rate more volatile.

F   Algorithm to solve the model

In the presence of aggregate shocks, market tightness depends on the distribution of positions and idiosyncratic productivity across firms, which is an infinite-dimensional object, as is common in models with heterogeneous agents. One method to solve the model numerically is to follow Krusell and Smith (1998) by approximating the infinite-dimensional state space using a finite number of aggregate state variables characterizing the distribution. I extend the Krusell and Smith (1998) algorithm by adding *one* specific additional state variable to capture omitted characteristics of the distribution: I add the residual in the law of motion for market tightness as a state variable. Thereby, I can increase the accuracy of the law of for market tightness significantly, and at the same time ensure that the market clears in each period.

Market tightness is an important variable, because it determines how costly it is for firms to fill their vacancies. When aggregating, the market clearing condition (14) must hold:

$$\int_{n_{i,-1}} \int_{x_i} \left(n^* \left(n_{i,-1}, x_i, s \right) - (1 - \lambda) n_{i,-1} \right)_+ dF_s \left(n_{i,-1}, x_i \right) = a \left(\theta \left(s \right) \right) u \left(s \right).$$

(F.1)

I denote by x_θ the residual in the law of motion for market tightness. The other aggregate state variables are aggregate productivity (z), volatility of idiosyncratic productivity (x_σ), the $\frac{1}{1-\alpha}$-moment of idiosyncratic firm productivity (m^x), and the aggregate number of positions at the beginning of the period (N_{-1}).

Note that the laws for z, x_σ, and m^x only depend on the exogenous processes chosen for aggregate productivity, idiosyncratic productivity, and the volatility of idiosyncratic productivity. Hence they are known (z, x_σ), or can be estimated (m^x) without solving the model.

For θ, x_θ, and N, I specify laws that are (log-)linear in all state variables:

$$\log \theta = \Psi_{\theta,0} + \Psi_{\theta,1} \log z + \Psi_{\theta,2} \log N_{-1} + \Psi_{\theta,3} x_\sigma + \Psi_{\theta,4} m^x + x_\theta, \qquad \text{(F.2)}$$

$$x_\theta = \Psi_{x_\theta,0} + \Psi_{x_\theta,1} \log z_{-1} + \Psi_{x_\theta,2} \log N_{-2} + \Psi_{x_\theta,3} x_{\sigma-1} + \Psi_{x_\theta,4} m^x_{-1} + \Psi_{x_\theta,5} x_{\theta-1},$$

$$\text{(F.3)}$$

$$\log N = \Psi_{N,0} + \Psi_{N,1} \log z_{-1} + \Psi_{N,2} \log N_{-1} + \Psi_{N,3} x_{\sigma-1} + \Psi_{N,4} m^x_{-1} + \Psi_{N,5} x_{\theta-1}.$$

$$\text{(F.4)}$$

Given these laws of motion (together with the known transition matrices for the aggregate shocks), one can find the optimal cutoffs in each state on the aggregate grid by iterating the marginal value function. From the cutoffs, one can calculate the number of positions each firm wants to fill, which depends on the idiosyncratic and aggregate state variables:

$$\left(n^* \left(n_{i,-1}, x_i, z, x_\sigma, m^x, N_{-1}, x_\theta \right) - (1 - \lambda) \, n_{i,-1} \right)_+ .$$

$$\text{(F.5)}$$

Then, I simulate the economy using the histogram method developed by Young (2010).[43] In the simulation, I impose the market clearing condition (F.1) using tthat all arguments except x_θ are predetermined in (F.5). I can solve for the unknown x_θ such that the market clearing condition (F.1) holds. Using the simulated values, the coefficients in the laws of motion (F.2), (F.3), and (F.4) can be estimated and updated. This procedure of solving the individual problem, simulating the economy, and updating the coefficients is repeated until the updated coefficients in one step are close to the ones estimated in the previous step.

---

[43]This means that insted of simulating a large number of firms, the cross-sectional distribution of positions and idiosyncratic productivity is characterized by a histogram over a fine grid. The advantage of this non-stochastic simulation method is that it avoids sampling error.